Stablecoin Architecture, Digital Dollarization, and the New Triffin Dilemma

Silvia Sgherri

George Washington University

This version: March 26, 2026

Abstract

Private stablecoins are not just payment technologies: they are monetary architectures that rewire global currency risk. Backed by sovereign debt and circulating outside any multilateral framework, they redistribute currency exposure across segmented international asset markets. This paper shows that the *design* of these instruments—not just their reserve quality—is the key policy problem. A three-country New Keynesian environment with portfolio-balance frictions and micro-founded liquidity services separates three design margins that are typically conflated: denomination, unit of account, and backing composition. The baseline is symmetric across countries in preferences, policy rules, and absorption of stablecoin positions across segmented asset markets. Under these conditions, each margin operates through a distinct equilibrium channel. Denomination determines portfolio-balance wedges and exchange-rate spillovers; the unit of account governs purchasing-power insurance through the liquidity channel; and backing composition shapes issuer's balance sheet valuations through bond-price and exchange-rate exposure. At the knife-edge benchmark, the three margins are orthogonal: symmetric denomination eliminates the stablecoin-induced component of all bilateral UIP premia; the unit-of-account basket affects real allocations and welfare without directly entering asset-market wedges; backingdenomination alignment insulates issuer balance sheets from exchange-rate valuation losses. Off the knife-edge, orthogonality breaks and welfare is primarily shaped by the unit-of-account margin, while denomination operates through second-order financial spillovers. Along the efficient-denomination locus, a sharp distributional conflict emerges, reflecting a structural trade-off between stabilization and insurance: the welfare-maximizing basket for the periphery conflicts irreconcilably with that of the core, and the conflict widens with adoption. Private stablecoins have thus recreated *the Triffin dilemma in digital form*, turning it from a question of supply into a problem of design. Denomination, unit of account, and backing composition are separate policy instruments, but their optimal design requires close coordination across them.

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

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

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 reconstructed the Triffin structure at scale, outside any multilateral governance framework, and at 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 oversight (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 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 'digital dollarization'—the substitution of domestic currencies with a concentrated dollar-denominated stablecoin system—is working exactly as designed.

The difficulty is that it generates externalities that no private issuer governs and no existing institution is mandated to resolve. Peripheral risk-off shocks transmit into advanced-economy exchange rates through stablecoin portfolio-balance channels, creating spillovers that are entirely absent in the no-stablecoin world. The conditions that maximize peripheral welfare simultaneously generate welfare costs for the core, and the conflict intensifies with every unit of adoption. Stablecoins thus reconstruct the inherent tension of the Triffin dilemma in digital form, shifting it from a question of liquidity supply to one of instrument design.

A structural framework of stablecoin architecture is required to characterize this mechanism. To this aim, the paper develops a three-country New Keynesian model with micro-founded liquidity services and portfolio-balance frictions. A single global stablecoin circulates across all three economies. Stablecoin design operates along three distinct and identifiable marginsdenomination, unit of account, and backing composition. Conflating them, as the existing literature largely does, leads to a misidentification of the first-order policy problem.

The model baseline is symmetric across countries in preferences, and policy rules and features symmetric absorption of stablecoin positions across segmented asset markets. However, while Home (H) and Foreign (F) are identical advanced-economy currency areas, the Rest of the World (R) is a high-yield peripheral economy whose assets carry a risk premium.

The only asymmetry hence arises from risk exposure: this asymmetry captures the empirical distinction between advanced economies and the periphery without imposing differences in preferences or technology.

Under these conditions, each stablecoin design margin operates through a distinct equilibrium channel generating different international spillovers. The *denomination vector* θ^d governs how fluctuations in outstanding supply translate into net currency exposures across segmented asset markets, entering UIP conditions via portfolio-balance premia. The *unit-of-account basket* θ^u defines the price index used to compute real purchasing power, governing the stablecoins insurance value relative to domestic money. The *backing composition* θ^a specifies the asset portfolio held in reserve, determining how exchange rate movements affect issuer equity. These three margins enter disjoint equilibrium conditions at the benchmark.

At the knife-edge benchmark, the three margins are exactly orthogonal: symmetric denomination eliminates the stablecoin-induced component of all bilateral UIP premia; the unit-of-account basket affects real allocations and welfare without directly entering asset-market wedges; backingdenomination alignment insulates issuer balance sheets from exchange-rate valuation losses. These results are structural properties of the benchmark environment. Off the knife-edge, orthogonality breaks and the three margins interact: welfare is shaped primarily by the unit-of-account margin, while denomination operates through second-order financial spillovers that scale with currency imbalances.[1]

Welfare analysis reveals a sharp distributional conflict along the efficient-denomination locus: the welfare-maximizing basket for the periphery conflicts irreconcilably with that of the core, and the conflict widens with adoption. Advanced economies prefer a basket design tilted toward safe currencies, which stabilize domestic conditions, while the peripheral economy prefers a basket design that insures against domestic-currency depreciation. The world optimum lies strictly between these positions.[2] Private stablecoins have therefore recreated *the Triffin dilemma in digital form*, turning it from a question of supply into a problem of design.

These results confirm that stablecoins do not just move money across borders: they also redistribute currency risk across countries. In this perspective, denomination, unit of account, and backing composition are separate policy instruments that require close coordination, as their interaction governs both financial stability and welfare. The original Triffin

[1] While exact neutrality holds at the symmetric knife-edge, the underlying mechanism is locally robust: financial spillovers scale linearly with denomination imbalances, so small deviations from equal weights generate only second-order welfare effects relative to concentrated designs. The knife-edge is therefore best interpreted as the center of a locally robust region rather than a fragile point.

[2] The welfare conflict is a quantitative implication of the models structural asymmetry between stabilization and insurance: its strength and location depend on the persistence and magnitude of peripheral risk, the importance of liquidity services, and relative country weights.

dilemma arose because liquidity provision and currency credibility were tied to a single instrument. The digital version differs fundamentally: the design margins are distinct and, in principle, separable. But because the welfare objective is non-separable, the marginal value of each instrument depends on the level of the other. Institutional separation does not therefore eliminate the underlying trade-off; it merely redistributes it across policy domains. Optimal stablecoin design requires close coordination across margins, as their interaction governs both financial stability and welfare. Whether the political economy of the international monetary system can deliver it remains an open question.

Relation to the Literature

Three debates in the literature prepare the ground for this paper.

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). The contribution here 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 denomination-wedge mechanism that generates this channel builds directly on the micro-founded portfolio-balance literature (Gabaix and Maggiori, 2015a; Itskhoki and Mukhin, 2021; Camanho et al., 2022), which establishes how intermediary balance-sheet constraints and capital-flow frictions produce endogenous UIP deviations; the stablecoin denomination vector is, in that sense, a structural shifter of the same wedge those papers identify. The three-country structure 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., 2020a; Gopinath and Stein, 2021a; Gourinchas and Rey, 2007; Gourinchas et al., 2017). This strand asks who supplies the safe asset and at what cost to the issuer. The present paper asks what happens when that asset is supplied by a private intermediary whose design choices are made for profit rather than systemic stability. The answer inverts the welfare hierarchy: the reserve-quality debate, which focuses on backing composition, addresses a second-order lever. The unit-of-account

channel generates welfare variation orders of magnitude larger. The instrument 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, and has focused entirely on dollar-denominated designs, for which neutrality is impossible. The empirical premise that stablecoins circulate as dollar substitutes and therefore inherit the vehicle-currency pricing dynamics of the dollar system is grounded in the currency internationalization literature (Gopinath and Stein, 2021b; Maggiori et al., 2020b), which documents how dollar invoicing and portfolio home bias jointly sustain the dollar's unit-of-account dominance; a non-dollar stablecoin must displace precisely those frictions to escape the neutrality knife-edge. The contribution is a framework with a clear structural characterization: denomination neutrality holds at the knife-edge benchmark; the unit-of-account and denomination channels are structurally orthogonal there, not merely empirically separable; and the New Triffin Dilemma admits a formal welfare characterization that identifies the conflict, quantifies it, and locates it along the adoption trajectory. The dilemma is not merely metaphoricalit is a structural implication of the model.

The paper proceeds as follows. Section 2 develops the model. Section 3 derives the main propositions. Section 4 presents the quantitative analysis in six steps: the core contagion result, the transmission mechanism, the real consequences, the balance-sheet doom loop, the welfare architecture and the distributional conflict, and the robustness of the results. Section 5 develops the two-instrument regulatory framework. Section 6 concludes. Appendices contain the full log-linearized system, the Blanchard–Kahn conditions, the derivation of the consumption-equivalent welfare, and the proofs of the main propositions.

2 Model Environment

The environment consists of three economies—Home (H), Foreign (F), and Rest of the World (R)—sharing a symmetric calibration with nominal rigidities, Taylor-rule monetary policy, and segmented international bond markets. H and F are identical advanced economies with deep and liquid bond markets, while R is a peripheral economy whose assets carry a risk premium. The key friction is imperfect substitutability between sovereign bonds: when assets are perfect substitutes, denomination is irrelevant; when substitutability is imperfect, denomination mechanically reallocates risk across private balance sheets.

The theoretical model formalizing this environment embeds a privately-issued global stablecoin backed by sovereign bond portfolios into an otherwise standard three-country New Keynesian framework. The stablecoin affects the economy through three distinct microfounded channels: denomination enters uncovered interest parity (UIP) conditions via portfolio allocation; the unit of account enters household money demand; and backing governs the issuer's balance sheet and valuation. All equilibrium conditions are derived from explicit optimization or arbitrage problems. All variables are log-deviations from a symmetric non-stochastic steady state. Throughout, superscript $^{+1}$ denotes the one-period-ahead date: $x^{+1} \equiv x_{t+1}$ and $\mathbb{E}_t x^{+1} \equiv \mathbb{E}_t x_{t+1}$.

2.1 Firms, Supply Side, and Monetary Policy

Firms in country $j \in \{H, F, R\}$ are monopolistically competitive and produce differentiated varieties using linear technology $y_j = n_j$, where n_j is labor input. Final goods are CES aggregates of varieties with elasticity of substitution $\mu > 1$, implying a steady-state markup of $\mu/(\mu - 1)$. Competitive labor markets and the linear technology imply that log-linearized real marginal cost equals output: $mc_j = y_j$, so the NKPC slope maps output directly to inflation.

Domestic prices are set subject to Calvo frictions (Calvo, 1983): in each period, a fraction $1 - \alpha$ of firms can reset prices optimally, while a fraction α keeps prices unchanged. The optimal reset price solves a forward-looking profit maximization problem under demand constraints, yielding the standard New Keynesian Phillips Curve:

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

where π_{jd} denotes domestic price inflation in country j's own currency and the slope coefficient is:

$$\kappa = \frac{(1-\alpha)(1-\beta\alpha)}{\alpha} \cdot \frac{1}{\mu}. \tag{2}$$

Price rigidity and desired markups jointly determine the responsiveness of inflation to real activity. Under the baseline calibration ($\alpha = 0.75$, $\mu = 6$) this implies $\kappa = 0.125$; the welfare weight on inflation dispersion is $\kappa/\mu = 0.021$ (see Appendix C).

CPI inflation aggregates domestic and imported components under CES preferences. With total import share γ split symmetrically across trading partners, each partner's com-

ponent enters with weight $\gamma/2$:

$$\pi_H = (1-\gamma)\,\pi_{Hd} + \tfrac{\gamma}{2}(\pi_{Fd} + \dot{e}_{HF}) + \tfrac{\gamma}{2}(\pi_{Rd} + \dot{e}_{HR}), \tag{3}$$

$$\pi_F = (1-\gamma)\,\pi_{Fd} + \tfrac{\gamma}{2}(\pi_{Hd} - \dot{e}_{HF}) + \tfrac{\gamma}{2}(\pi_{Rd} + \dot{e}_{FR}), \tag{4}$$

$$\pi_R = (1-\gamma)\,\pi_{Rd} + \tfrac{\gamma}{2}(\pi_{Hd} - \dot{e}_{HR}) + \tfrac{\gamma}{2}(\pi_{Fd} - \dot{e}_{FR}). \tag{5}$$

Here $\dot{e}_{jk}$ denotes the log change of the nominal exchange rate e_{jk} (units of currency j per unit of currency k), so that $\dot{e}_{jk} > 0$ corresponds to a depreciation of j relative to k.

Each central bank follows an inertial Taylor rule with smoothing parameter $\rho_i \in (0,1)$, Taylor-principle coefficient $\phi_\pi > 1$, and output coefficient $\phi_y > 0$:

$$i_j = \rho_i\, i_j^{-1} + (1-\rho_i)(\phi_\pi \pi_j + \phi_y\, y_j)\,. \tag{6}$$

2.2 Households and Liquidity Demand

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

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_{j,t}^{1-\sigma}}{1-\sigma} - \frac{\varphi_l}{2}(C_{j,t} - L_{j,t})^2 \right], \tag{7}$$

where the second term captures the utility value of liquidity services.

Liquidity services $L_{j,t}$ aggregate domestic money $MB_{j,t}$ and stablecoin holdings $S_{j,t}$. Under the limit of infinite elasticity of substitution between money and stablecoin, the CES aggregator log-linearizes to the weighted arithmetic mean:

$$l_j = (1-\omega)\, mb_j + \omega\, s_j, \tag{8}$$

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

$$(\sigma + \varphi_l)\, c_j = (\sigma + \varphi_l)\, \mathbb{E}_t c_j^{+1} - \varphi_l\, \mathbb{E}_t(l_j^{+1} - l_j) - (i_j - \mathbb{E}_t \pi_j^{+1}), \tag{9}$$

so expected changes in liquidity shift the IS curve directly via $\varphi_l\, \mathbb{E}_t(l_j^{+1} - l_j)$. Standard intratemporal money demand follows from the household first-order condition:

$$mb_j = -\eta i_j. \tag{10}$$

The relative demand for stablecoins versus money is governed by the intratemporal optimality condition equating the marginal rate of substitution between liquid assets to their

relative opportunity cost. The relative cost of holding stablecoins versus money equals the inflation differential $\pi_j^{\text{sc}} - \pi_j$: the stablecoin provides a lower real return when its basket inflates faster than domestic prices. The resulting log-linear condition is:

$$mb_j - s_j = \eta(\pi_j^{\text{sc}} - \pi_j), \tag{11}$$

where π_j^{sc} is stablecoin basket inflation in country j's own numeraire. When $\pi_j^{\text{sc}} > \pi_j$ households shift toward domestic money; when $\pi_j^{\text{sc}} < \pi_j$ they shift toward stablecoins. This is the primary channel through which the unit-of-account vector affects the real economy.

2.3 The Stablecoin: Design Parameters and Equilibrium Channels

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

Denomination and the UIP wedge. International bond markets are segmented: households face imperfect substitutability across currency denominations, modelled as a quadratic portfolio adjustment cost $\frac{\psi_b}{2} b_{jk}^2$ on each bilateral NFA b_{jk}, additively separable across pairs. Minimizing this cost delivers a uniform portfolio-balance coefficient ψ_b for all bilaterals under the standard symmetric calibration.

Stablecoin issuance shifts the effective NFA position by introducing an additional currency exposure into segmented asset markets. In portfolio-balance environments, the relevant object is the net currency position that must be absorbed by price-sensitive balance sheets in each bilateral market.

In the present framework, this exposure is characterized at the level of aggregate stablecoin supply. This corresponds to a benchmark environment in which (i) stablecoin liabilities represent a common claim on a globally integrated intermediary, (ii) the associated currency exposure is transmitted uniformly across segmented markets, and (iii) portfolio-balance coefficients are symmetric across countries and bilaterals. Under these conditions, one unit of stablecoin supply is equivalent, in H numeraire, to a synthetic currency portfolio with denomination weights $(\theta_H^d, \theta_F^d, \theta_R^d)$.

Let $\delta_{jk} \equiv \theta_j^d - \theta_k^d$ denote the bilateral denomination imbalance. Aggregate issuance $\bar{s} \cdot s$ then introduces the effective bilateral exposure

$$\tilde{b}_{jk} = b_{jk} + \delta_{jk}\, \bar{s} \cdot s.$$

This formulation is a special case of a more general mapping in which the effective currency imbalance depends jointly on denomination and on how stablecoin positions are absorbed across countries and markets. The knife-edge neutrality result derived below obtains exactly under the benchmark conditions stated above. Departures from these conditions affect the mapping between denomination and market-level imbalances but do not alter the central role of denomination in governing the direction and magnitude of portfolio-balance distortions.

At steady-state adoption $\bar{s} = \kappa_S \omega$ (with $\kappa_S = 1$, so $\bar{s} = \omega$), aggregate issuance $\bar{s} \cdot s$ introduces a net jk imbalance $\delta_{jk}\bar{s} \cdot s$ into bilateral asset markets. The *stablecoin-induced effective bilateral exposure* is therefore:

$$\tilde{b}_{jk} = b_{jk} + \delta_{jk}\bar{s} \cdot s, \tag{12}$$

where b_{jk} is the household bilateral NFA in the absence of stablecoins. The second term captures the denomination-induced currency imbalance proportional to aggregate supply s. Under the benchmark conditions described abovenamely, a common intermediary structure, symmetric portfolio-balance coefficients, and uniform transmission of currency exposure across segmented marketsthis exposure is equivalent to that generated by a synthetic currency portfolio with weights $(\theta^d_H, \theta^d_F, \theta^d_R)$, and therefore depends on denomination rather than on the geographic distribution of holdings. [3]

Applying the portfolio-balance condition to $\tilde{b}_{jk}$ yields the stablecoin-induced UIP wedges:

$$\psi^{\text{sc}}_{HF} = \psi_b \, \delta_{HF} \, \bar{s}(s - \zeta_E E^{\text{sc}}), \tag{13}$$

$$\psi^{\text{sc}}_{HR} = \psi_b \, \delta_{HR} \, \bar{s}(s - \zeta_E E^{\text{sc}}), \tag{14}$$

$$\psi^{\text{sc}}_{FR} = \psi^{\text{sc}}_{HR} - \psi^{\text{sc}}_{HF} = \psi_b \, \delta_{FR} \, \bar{s}(s - \zeta_E E^{\text{sc}}), \tag{15}$$

where $\zeta_E \geq 0$ captures a secondary balance-sheet effect: depleted issuer equity contracts the effective stablecoin supply available for portfolio rebalancing. The F-R wedge (15) is not an independent condition; it follows from the triangular no-arbitrage identity $e_{FR} = e_{HR} - e_{HF}$.

Setting $\theta^d_H = \theta^d_F = \theta^d_R = 1/3$ zeros all three imbalances simultaneously and collapses $\tilde{b}_{jk}$ back to b_{jk} for every bilateral. This is the economic content of Proposition 1: *symmetric denomination*—equal weights on all three currencies, also referred to throughout as the *knife-edge*—creates a stablecoin with zero net currency imbalance in every bilateral market, leaving the portfolio-balance channel undisturbed at any adoption scale $\bar{s}$.

[3] More generally, relaxing these conditions introduces a dependence of the effective imbalance on the distribution of holdings, but denomination remains the sufficient statistic for the direction of currency exposure.

Unit of account and money demand. Let the stablecoin price index be the weighted average of constituent prices in H numeraire. Log-linearizing, stablecoin basket inflation in H numeraire is:

$$\pi^{\text{sc}} = \theta^u_H \pi_{Hd} + \theta^u_F(\pi_{Fd} + \dot{e}_{HF}) + \theta^u_R(\pi_{Rd} + \dot{e}_{HR}). \tag{16}$$

Basket inflation in each country's own numeraire is obtained by adjusting for the bilateral rate change:

$$\pi^{\text{sc}}_H = \pi^{\text{sc}}, \qquad \pi^{\text{sc}}_F = \pi^{\text{sc}} - \dot{e}_{HF}, \qquad \pi^{\text{sc}}_R = \pi^{\text{sc}} - \dot{e}_{HR}. \tag{17}$$

The unit-of-account vector hence governs how foreign inflation and exchange-rate movements map into the stablecoin's purchasing power in each country. Through the arbitrage condition (11), country-specific deviations between π^{sc}_j and domestic CPI shift household portfolios between money and stablecoins, feeding directly into the IS curve via the liquidity term in (9).

By inspection of equations (13)–(15), the unit-of-account vector θ^u does not appear in any UIP wedge: θ^u affects relative prices and money demand, but not portfolio-balance premia. Denomination and unit-of-account channels hence enter disjoint equilibrium equations.

At the knife-edge (where $\delta_{jk} = 0$ for all jk), this orthogonality is exact: aggregate demand s still responds to π^{sc} through money demand (11), but the resulting movement in s introduces no net currency exposure because $\psi^{\text{sc}}_{jk} = \psi_b\, \delta_{jk}\, \bar{s}(\cdot) \equiv 0$ regardless of s. Off the knife-edge an indirect channel through s exists, but its magnitude is proportional to δ_{jk} and represents a quantitatively negligible leakage from the unit-of-account channel into the financial block relative to its primary transmission through money demand, as confirmed in Table 2 (Proposition 4).

Stablecoin aggregation. Aggregate stablecoin supply reflects the distribution of stablecoin holdings:

$$s = \lambda_s s_H + \tfrac{1-\lambda_s}{2} s_F + \tfrac{1-\lambda_s}{2} s_R. \tag{18}$$

The symmetric baseline $\lambda_s = 1/3$ assigns equal weight to all three countries.

2.4 Exchange Rate Determination and Financial Structure

Bilateral exchange rates are determined by no-arbitrage conditions in international asset markets. We impose triangular arbitrage so that two bilateral exchange rates—e_{HF} (H per F) and e_{HR} (H per R)—are independently determined by uncovered interest parity conditions

with portfolio-balance premia (Gabaix and Maggiori, 2015b; Obstfeld et al., 2009):

$$e_{HF} = \mathbb{E}_t e_{HF}^{+1} + i_H - i_F - \psi_{HF}, \tag{19}$$

$$e_{HR} = \mathbb{E}_t e_{HR}^{+1} + i_H - i_R - \psi_{HR}. \tag{20}$$

The e_{FR} rate satisfies triangular no-arbitrage and is derived residually: $e_{FR} = e_{HR} - e_{HF}$.

The premia ψ_{jk} arise from imperfect asset substitutability and decompose into an NFA term, the stablecoin wedge defined in (13)–(14), an equity-feedback term, and a risk shock:

$$\psi_{HF} = \psi_b \, b_{HF} + \psi_{HF}^{\text{sc}} - \chi_E \, E^{\text{sc}}, \tag{21}$$

$$\psi_{HR} = \psi_b \, b_{HR} + \psi_{HR}^{\text{sc}} - \chi_E \, E^{\text{sc}} + \nu. \tag{22}$$

The F-R premium is derived residually as $\psi_{FR} = \psi_{HR} - \psi_{HF}$, preserving triangular consistency.

The global risk-off shock ν enters ψ_{HR} only and follows an AR(1) process:

$$\nu_t = \rho_\nu \, \nu_{t-1} + \varepsilon_t^\nu, \qquad \varepsilon_t^\nu \sim \text{iid}(0, \sigma_\nu^2). \tag{23}$$

Under triangular no-arbitrage, the shock ν transmits to the F-R premium residually through $\psi_{FR} = \psi_{HR} - \psi_{HF}$: when ν raises the cost of R borrowing against H, it raises it against F by the same amount, since ψ_{HF} is undisturbed in the absence of stablecoin-induced spillovers. The H-F bilateral is therefore insulated from the shock by construction in the no-stablecoin world, and only by architecture in the stablecoin world.[4]

The equity-feedback term $-\chi_E E^{\text{sc}}$ arises from the issuer's role as a financial intermediary and governs the extent to which issuer equity dampens or amplifies the overall premium: when equity is depleted, the issuer's effective risk-bearing capacity falls, raising the premium required to clear the market. This generates what we term a *doom-loop amplification*: a persistent feedback mechanism in which valuation losses reduce equity, lower equity raises premia, and higher premia induce further valuation losses. In the linearized system, this operates through persistence rather than through a discrete collapse threshold.

[4] In the triangular world, H-F insulation is not automatic even without stablecoins: it is a consequence of the arbitrage structure. With stablecoins and concentrated denomination, that insulation breaks, as established in Section 2.3.

2.5 Issuer Balance Sheet and Equity Dynamics

The issuer solves a standard portfolio problem subject to its backing constraints. The nominal return on its backing portfolio in H numeraire is:

$$r^A = \theta^a_H \, i_H + \theta^a_F (i_F + \mathbb{E}_t \dot{e}^{+1}_{HF}) + \theta^a_R (i_R + \mathbb{E}_t \dot{e}^{+1}_{HR}). \tag{24}$$

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

$$E^{\text{sc}}_t = (\beta^{-1} - \phi_{\text{div}}) E^{\text{sc}}_{t-1} + \pi_{\text{prof}} + \text{val}_q + \text{val}_{fx}, \tag{25}$$

where $\text{val}_q = \bar{s} \sum_j \theta^a_j \Delta q_j$ is the bond-price valuation term, with $q_{j,t} = -i_{j,t}/(1-\beta)$ the log-linearized real bond price. The FX valuation term,

$$\text{val}_{fx} = \bar{s}\big[(\theta^a_F - \theta^d_F)\dot{e}_{HF} + (\theta^a_R - \theta^d_R)\dot{e}_{HR}\big], \tag{26}$$

follows from the currency mismatch between assets and liabilities. Under backing-denomination alignment $\theta^a_j = \theta^d_j$ for all j, equation (26) vanishes identically and equity is exchange-rate insulated. Under misalignment with $\theta^a_R > \theta^d_R$—the issuer overweights R assets relative to R-denominated liabilities—R depreciation ($\dot{e}_{HR} < 0$) generates $\text{val}_{fx} < 0$: equity is destroyed procyclically.

At the baseline $\phi_{\text{div}} = 1$, the autoregressive coefficient on E^{sc}_{t-1} equals $\beta^{-1} - 1 \approx 0.01$ at quarterly frequency, making equity approximately a static accounting identity with no independent causal role. Buffers are activated in Section 4.5 by switching to $\phi_{\text{div}} = 0.5$, $\chi_E = 0.025$, $\zeta_E = 0.05$.

Dividends are rebated to households in proportion to their stablecoin holdings. Since λ_s in (18) is H's share of aggregate stablecoin supply, H households receive fraction λ_s of total dividends: $T_H = \lambda_s \, \phi_{\text{div}} E^{\text{sc}}_{t-1}$, and similarly for F and R with weights $(1-\lambda_s)/2$ each.

2.6 External Balance and Goods Market Clearing

Bilateral terms of trade evolve as:

$$\dot{\text{tot}}_{HF} = \dot{e}_{HF} + \pi_{Fd} - \pi_{Hd}, \tag{27}$$

$$\dot{\text{tot}}_{HR} = \dot{e}_{HR} + \pi_{Rd} - \pi_{Hd}, \tag{28}$$

$$\dot{\text{tot}}_{FR} = \dot{e}_{FR} + \pi_{Rd} - \pi_{Fd}. \tag{29}$$

Only two are independent: $\text{tot}_{FR} = \text{tot}_{HR} - \text{tot}_{HF}$ by construction. 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 yields $y_j = c_j + nx_j - T_j$ 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 D. □

The result is structural, not a calibration artifact. 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.

Corollary 1 (Stablecoin Monetary Neutrality by Design)**.** *A CBDC issued by country j achieves monetary neutrality with respect to all bilateral UIP premia by the nature of sovereign issuance: as a direct liability of the sovereign monetary authority it is part of the monetary base, so the portfolio-balance wedge formula (13)–(15) does not apply—there is no private intermediary creating a currency mismatch in international asset markets. A privately issued stablecoin achieves the same neutrality outcome* if and only if *it adopts the knife-edge denomination $\theta_H^d = \theta_F^d = \theta_R^d = 1/3$, zeroing the wedge by design rather than by institutional definition. Unlike CBDC neutrality, which is a structural consequence of sovereign issuance, stablecoin neutrality is a* design choice*—one that can be mandated by regulation without restricting the stablecoin's unit-of-account properties or its role as peripheral currency insurance.*

Proof. See Appendix D. □

3.2 Cross-Market Contagion and Its Elimination

Proposition 2 (Cross-Market Insulation)**.** *Under symmetric diversified denomination, the risk-off shock ν entering ψ_{HR} has no first-order effect on ψ_{HF} or e_{HF}, even through its induced changes in aggregate stablecoin demand s. Under concentrated denomination ($\theta_H^d >$*

θ_F^d, *so* $\delta_{HF} > 0$*), changes in* s *enter* ψ_{HF}^{sc} *with coefficient* $\psi_b\,\delta_{HF}\,\bar{s} > 0$*, raising* ψ_{HF} *and transmitting the peripheral shock directly into the H-F market.*

Proof. See Appendix D. □

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: 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 realizations. 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 D. □

Equity feedback and amplification. When equity feedback is active ($\chi_E > 0$), issuer equity enters risk premia linearly through $-\chi_E E^{sc}$. Negative equity realizations therefore amplify premia, while positive equity dampens them. Under backing–denomination misalignment, exchange-rate movements generate procyclical valuation losses (via val^{fx}) that can drive equity negative on impact, creating a state-contingent amplification mechanism.

This effect is a property of the linearized system, not a nonlinear regime switch. The locus $E^{sc} = 0$ is not a binding constraint but a reference point at which the contribution of equity to premia changes sign continuously. The amplification mechanism described here therefore operates within a local approximation and should be interpreted as a propagation channel rather than as a statement about default or market breakdown.

3.4 Unit-of-Account Orthogonality

Proposition 4 (Unit-of-Account Orthogonality)**.** *[(i)]*

1. Exact orthogonality at the knife-edge. *Under symmetric diversified denomination (*$\theta_H^d = \theta_F^d = \theta_R^d = 1/3$*), the unit-of-account basket* θ^u *has no effect—direct or indirect—on any bilateral UIP premium, any exchange rate, or* val_{fx}*, at first order.*

2. Conditional orthogonality off the knife-edge. *For any fixed denomination vector θ^d with heterogeneous weights ($\theta^d_j \neq \theta^d_k$ for some j,k), θ^u affects ψ^{sc}_{jk} only through the indirect path $\theta^u \to \pi^{\text{sc}} \to s \to \psi^{\text{sc}}_{jk}$, with slope $\psi_b(\theta^d_j - \theta^d_k)\bar{s} \cdot \partial s/\partial\theta^u$. This indirect effect represents a quantitatively negligible leakage from the unit-of-account channel into the financial block relative to its primary transmission through money demand: the welfare hierarchy of Table 2 confirms that varying θ^u along the simplex edge while holding θ^d fixed generates welfare variation orders of magnitude larger than varying θ^d alone, consistent with the UoA channel operating primarily through money demand rather than through the premium.*[5]

In both cases, θ^u does not enter val_{fx} or the backing channel at any order.

Proof. See Appendix D. □

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 $\psi^{\text{sc}}_{HF} = \psi^{\text{sc}}_{HR} = \psi^{\text{sc}}_{FR} = 0$ and $\text{val}_{fx} = 0$ simultaneously for all exchange rate realizations.*

Proof. Direct from equations (13)–(15) and (26): $\psi^{\text{sc}}_{jk} = 0$ for all jk requires $\delta_{jk} = 0$ for all jk, i.e. $\theta^d_H = \theta^d_F = \theta^d_R = 1/3$; $\text{val}_{fx} = 0$ for all exchange rate realizations requires $\theta^a_j = \theta^d_j$ for all j. □

Perturbations around the knife-edge. Let $\theta^d = (1/3 + \epsilon_H,\ 1/3 + \epsilon_F,\ 1/3 + \epsilon_R)$ with $\epsilon_H + \epsilon_F + \epsilon_R = 0$ and define bilateral deviations $\delta_{jk} \equiv \theta^d_j - \theta^d_k = \epsilon_j - \epsilon_k$. For small deviations, the stablecoin-induced UIP wedges satisfy

$$\psi^{sc}_{jk} = \psi_b\, \delta_{jk}\, \bar{s}\, (s - \zeta_E E^{sc}),$$

so that the sensitivity of premia to denomination imbalances is linear in δ_{jk}.

[5] A stronger sufficient condition for exact independence of ψ^{sc}_{jk} with respect to θ^u at general θ^d would be inelastic stablecoin supply ($\partial s/\partial\theta^u = 0$) or zero money-stablecoin substitutability ($\eta = 0$). Neither is economically attractive. The quantitative evidence in Table 2 establishes that the indirect leakage into the financial block is negligible relative to the unit-of-account channel's primary transmission through money demand, even where strict theoretical separability does not hold.

It follows immediately that the exact orthogonality result at the knife-edge ($\delta_{jk} = 0$ for all jk) extends locally: deviations from neutrality scale proportionally with the magnitude of denomination imbalance. The indirect leakage from the unit-of-account channel into the financial block operates through the composite mapping

$$\theta^u \longrightarrow \pi^{sc} \longrightarrow s \longrightarrow \psi^{sc}_{jk},$$

with slope proportional to $\psi_b \, \delta_{jk} \, \bar{s} \cdot \partial s / \partial \theta^u$. Hence, for small $\|\epsilon\|$, the deviation from orthogonality is first-order in δ_{jk} and inherits the magnitude of equilibrium demand responses.

This characterization implies that the knife-edge is not a fragile point but the center of a locally robust region in which financial spillovers remain quantitatively small as long as denomination imbalances are limited.

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 design. The welfare analysis of Section 4.6 will show that the UoA dimension generates welfare variation orders 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-maximizing basket for R households is incompatible with the welfare-maximizing basket for the advanced-economy core. No design choice resolves this conflict from within. It requires governance from without.

4 Quantitative Analysis

The quantitative analysis proceeds 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 uniform portfolio-balance coefficient $\psi_b = 0.60$ across all bilaterals reflects the symmetric variance-covariance assumption required by the denomination-neutrality microfoundation; asymmetry between safe and risky

bilaterals enters exclusively through the shock structure. The portfolio-balance coefficient 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, 2015b).

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.

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; the 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 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 (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 generates 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 stablecoin holdings approximately three times the initial shock magnitude under Regime 3, and roughly 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, $\delta_{HF} = 1$ and the hump translates directly into ψ_{HF}. Under the knife-edge $\delta_{HF} = 0$ and no premium results, even though the aggregate hump in s is present.

The basket inflation panel π^{sc} (bottom-right) 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 pure H UoA, $\pi^{sc} \approx \pi_{Hd}$ and barely moves. The contrast is stark: the regime with violent basket inflation (Regime 4) is exactly the regime with zero ψ_{HF} response, and the regimes with stable basket inflation (Regimes 2 and 3) 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. The three-row decomposition thus identifies a clear hierarchy: the liquidity and consumption channels are architecture-sensitive; the external adjustment channel is architecture-neutral. This hierarchy is precisely what the welfare analysis in Section 4.6 will formalize.

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 (Propositions 3 and that the doom loop is a misalignment phenomenon that correctly aligned architecture prevents entirely.

The phase diagram in Figure 4 illustrates the equity-feedback amplification mechanism

with full geometric clarity. In the left panel ($\chi_E = 0$, no active feedback), the three regime trajectories reveal the architecture mechanism in its purest form. Regime 4 (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. Regime 2 (aligned H-backing) makes a much tighter loop near the origin, also in the stabilization zone. Regime 3 (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 $t = 0$ and the entire subsequent orbit lies in the red zone. 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. For Regime 4 (green dashed), activating feedback expands the stabilization loop: positive equity now dampens premia. For Regime 2 (blue dashed), the effect is small—the loop barely changes because aligned backing keeps equity positive. For Regime 3 (red dashed), the effect reverses and amplifies. 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. The paradox is resolved by the linear dynamics of the equity law of motion. Activating buffers reduces the payout rate from $\phi_{\text{div}} = 1$ to $\phi_{\text{div}} = 0.5$, which increases the autoregressive coefficient on E^{sc}_{t-1} in equation (25) from $(\beta^{-1} - 1) \approx 0$ to $(\beta^{-1} - 0.5) \approx 0.5$. Under Regime 3, the initial val_{fx} destruction on impact drives E^{sc}_t sharply negative. With full payout, this negative equity is distributed immediately and dissipates within one period. With partial retention, the same initial destruction persists for many periods, generating a larger discounted sum of squared equity deviations. When equity feedback is also active ($\chi_E > 0$), the persistent negative equity enters the risk premia through $-\chi_E E^{\text{sc}}_t$, depressing exchange rates in subsequent periods and triggering further val_{fx} losses through the misaligned backing. The increase in equity volatility under buffers is therefore not a levels story about a larger FX-exposed balance sheet: it is a persistence and feedback

story. This distinction is critical. In a linear approximation, buffers do not attenuate losses by scaling the balance sheet; they modify the persistence of equity dynamics. Larger buffers do not mechanically absorb shocks. Instead, partial retention prolongs the effect of a given valuation shock, and when feedback is active, this persistence becomes a source of amplification rather than stabilization. The right panel resolves the paradox. 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 (center) 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. 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. 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 -5.4×10^{-3} under Regime 3—losses that are three to four times larger than 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 with a recovery of approximately $+0.6 \times 10^{-3}$, leaving a net R welfare loss of -4.8×10^{-3}—more than 80 percent 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 center and right panels side by side produces the distributional verdict: the architecture most beneficial to the safe bloc (Regime 2, positive center 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 zeros out both bars simultaneously, eliminating both the safe-bloc gain and the R-country loss in one step.

Interpretation. The balance-sheet dynamics described in this section are derived from a first-order log-linear approximation. As such, they capture amplification around the steady state but abstract from nonlinear features such as default, liquidation constraints, or endogenous market illiquidity. The analysis should therefore be interpreted as characterizing local propagation mechanisms rather than global stability properties. A fully nonlinear treatment would be required to determine whether equity depletion leads to discontinuous outcomes such as issuer failure or market freezes. The results here instead show that, even absent such nonlinearities, misaligned architecture generates amplification forces that capital buffers alone cannot neutralize.

Remark. The linear framework isolates the amplification mechanism cleanly; introducing nonlinear constraints would likely strengthen, rather than overturn, the qualitative asymmetry between aligned and misaligned architectures.

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), we report 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 findings.

The welfare hierarchy. The dominant result is the unit-of-account margin's supremacy. Along the HR edge—which trades H-currency weight against R-currency weight with F weight set to zero—the UoA welfare range is 6.02×10^{-3} in Regime 2, 6.03×10^{-3} in Regime 3, and 4.76×10^{-3} in Regime 4. The corresponding denomination ranges are 0.91, 0.20, and

0.03×10^{-3}, yielding U/D ratios of 6.6, 29.7, and 193 respectively. The backing margin lies between them. The hierarchy is unambiguous: $\text{Range}^U \gg \text{Range}^A \gg \text{Range}^D$.

The mechanism is transparent. The denomination channel works through portfolio-balance wedges and reaches households indirectly via exchange-rate volatility and output gaps. The unit-of-account channel, by contrast, hits households directly in the utility function: when R depreciates and the stablecoin basket is safe-bloc 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 basket composition rather than currency substitution.

The knife-edge row provides a striking corollary. On the HF edge, the denomination range collapses to numerically zero; on the HR edge it is just 0.03×10^{-3}. The UoA range on the HR edge remains 4.76×10^{-3}, fully 79 percent of its Regime 2 value. The knife-edge eliminates denomination-channel welfare effects while leaving the UoA welfare gradient largely intact. *At the knife-edge, denomination regulation and UoA governance are welfare-independent policy instruments*—a conditional separability that breaks down off the knife-edge but is quantitatively robust in its neighborhood.

Distance from the knife-edge. To assess the quantitative robustness of the knife-edge prescription, we evaluate deviations from symmetric denomination. Consider a perturbation along the HR edge with $\theta^d = (1/3+\Delta,\ 1/3,\ 1/3-\Delta)$ so that $\delta_{HR} = 2\Delta$. Moving from $\Delta = 0$ to $\Delta = 0.04$ (i.e., $\theta^d_H = 0.37$) generates a strictly proportional increase in the stablecoin-induced UIP wedge and associated welfare losses.

Quantitatively, the resulting changes in welfare remain small relative to the gap between concentrated and diversified regimes: the welfare gradient in the neighborhood of the knife-edge is shallow, while the gradient toward concentration is steep. This asymmetry reflects the linear scaling of financial spillovers in δ_{jk} combined with the second-order nature of denomination effects relative to the unit-of-account channel.

The knife-edge should therefore be interpreted not as a point requirement but as the center of a tolerance region within which welfare losses remain negligible relative to those generated by concentrated designs.

Welfare non-separability. The conditional separability of instruments at the knife-edge does not extend to the welfare objective. Table 5 (Appendix C) demonstrates this quantitatively: it is the numerical counterpart to the tilted contour geometry in Figure 7.

Three results follow directly from comparing the two panels. First, the unit-of-account

margin dominates conditionally: at the knife-edge, $\Delta CE_R^U = 3.23 \times 10^{-3}$ versus $\Delta CE_R^D = 0.37 \times 10^{-3}$, a ratio of 8.8×. Second, this dominance is robust across denomination choices: ΔCE_R^U varies by only 20 percent over the full denomination range (Panel A), whereas ΔCE_R^D varies by 98 percent across the unit-of-account range (Panel B). Third, the two instruments are welfare substitutes for R: as $\bar{x}_u$ approaches R's preferred unit of account, the marginal welfare value of denomination collapses—ΔCE_R^D falls from 0.369×10^{-3} to 0.007×10^{-3} as $\bar{x}_u$ rises from 1/3 to 0.87. Strikingly, when the unit of account excludes R entirely ($\bar{x}_u = 0$), R's conditional optimum over denomination lies at the H corner rather than the R corner. This is because, at $\bar{x}_u = 0$, full R-denomination would generate substantial macroeconomic volatility through the UIP wedge. Although denomination provides insurance against R depreciation, this insurance operates through a distorted financial channel that amplifies fluctuations in consumption and the output gap. In equilibrium, the distortion dominates the benefit, so even R households do not choose R-denominated stablecoins on that slice—reflecting the underlying insurance-distortion trade-off in the model.

Together, these findings confirm that UoA is the primary welfare instrument while denomination, though essential for financial stability, is the secondary welfare lever. Crucially, these results establish that the two margins are substitutes locally—the marginal value of each declines as the other becomes aligned—but globally non-separable, so the optimum is jointly determined.

Proposition 5 (Welfare Non-Separability)**.** *The welfare objective $CE_j(x_d, x_u)$ is non-separable in denomination and unit-of-account weights. In particular, the optimal unit-of-account choice depends on denomination, and the optimal denomination choice depends on the unit of account. Consequently, the world optimum (x_d^*, x_u^*) requires joint optimization of the two margins and cannot be obtained by sequential optimization.*

Proof. See Appendix D. □

Figure 7 maps the joint variation of denomination and unit-of-account across the full two-dimensional design space. The welfare-loss contours are nearly horizontal: welfare loss is steep in the unit-of-account dimension and comparatively flat in the denomination dimension. The white star at $(x_d^*, x_u^*) \approx (0.12, 0.34)$ marks the unconstrained two-dimensional optimum. The non-axis-aligned contours confirm Proposition 5 geometrically: were the two instruments welfare-separable, the contours would be axis-aligned. They are not.

The New Triffin Dilemma The argmax columns in Table 2 expose the distributional conflict over the unit-of-account basket. Along the HF-R edge, the advanced-economy safe bloc prefers a low R-currency weight, while the peripheral economy prefers a high one; the world

optimum lies strictly between them. In the baseline knife-edge calibration, the constrained one-dimensional search reveals that the advanced-economy safe bloc prefers $x^*_{HF} \approx 0.09$–0.13 while R prefers the opposite, $x^*_R \approx 0.87$–0.89. The world optimum $x^*_W \approx 0.33$–0.37 lies strictly between these positions and satisfies neither. The corresponding two-dimensional search confirms the same ordering: the HF optimum lies at $(0.00, 0.09)$, the R optimum at $(0.19, 0.87)$, and the world optimum at $(0.12, 0.34)$.[6] No single basket simultaneously maximizes peripheral and advanced-economy welfare. The welfare conflict $\text{Loss}_j(\omega)$ defined in equations (30)–(31) is strictly increasing in adoption ω: the distributional conflict intensifies as stablecoin adoption grows. That is the new Triffin dilemma in stablecoin design space.

Analytical characterization. The ordering $x^*_{HF} < x^*_R$ is not an algebraic identity of the model in the way denomination neutrality is. It is a quantitative implication of the models structural asymmetry between advanced-economy stabilization and peripheral insurance demand. Along the HR edge, increasing the R-currency weight in the unit-of-account basket raises the extent to which the stablecoin insures R households against domestic-currency depreciation, but it also increases the volatility of stablecoin purchasing power relative to the CPI of H and F. Formally, the sign of the welfare gradient $\partial CE_j/\partial x_u$ reflects the balance between an insurance effect (operating through the alignment of stablecoin inflation with domestic marginal utility in R) and a stabilization cost (operating through money-demand distortions and the IS equation in H and F). The peripheral optimum exceeds the advanced-economy optimum whenever the marginal insurance gain in R dominates the induced domestic stabilization cost there, while the corresponding marginal effect is negative for the safe bloc. This condition is more likely to hold when: (i) the risk-premium shock affecting R is large; (ii) its persistence ρ_ν is high, so that insurance is valuable in discounted terms; (iii) the liquidity wedge is quantitatively important, so that deviations between stablecoin inflation and domestic CPI carry substantial welfare weight; and (iv) the peripheral economy receives sufficient welfare weight, either through country size or through the severity of its exposure. These primitives characterize the region of parameter space in which the conflict emerges, even though they do not imply a closed-form global ordering.

The coordination gain. Figure 8 renders the Triffin dilemma in its starkest form. Define the *welfare conflict* for each bloc as the CE welfare cost of being forced to the opposing

[6] The HR-edge constrained search (Table 2) gives $x^*_W = 0.35$ because it fixes $x_d = 1/3$; the unconstrained two-dimensional optimum gives $x^*_u = 0.34$ at $x^*_d = 0.12$ (Figure 7). The difference is economically negligible but reflects the complementarity between a slightly lower denomination R-weight and a slightly lower UoA R-weight at the global welfare maximum.

bloc's two-dimensional optimal basket:

$$\text{Loss}_R(\omega) = CE_R(x^*_R;\, \omega) - CE_R(x^*_{HF};\, \omega) \geq 0, \tag{30}$$

$$\text{Loss}_{HF}(\omega) = CE_{HF}(x^*_{HF};\, \omega) - CE_{HF}(x^*_R;\, \omega) \geq 0, \tag{31}$$

where $x^*_{HF} \approx (0.00, 0.09)$ and $x^*_R \approx (0.19, 0.87)$ are the full two-dimensional bloc optima from Table 4 Panel A. The *coordinator gain* is the welfare each bloc recovers when a benevolent planner imposes the world-welfare-optimal basket $x^*_W \approx (0.12, 0.34)$ rather than the opposing bloc's preferred basket:

$$\text{CG}_R(\omega) = CE_R(x^*_W;\, \omega) - CE_R(x^*_{HF};\, \omega) \geq 0, \tag{32}$$

$$\text{CG}_{HF}(\omega) = CE_{HF}(x^*_W;\, \omega) - CE_{HF}(x^*_R;\, \omega) \geq 0. \tag{33}$$

The residual loss $\text{Loss}_j - \text{CG}_j \geq 0$ is the conflict the coordinator cannot resolve.

The left panel of Figure 8 plots CE_R, CE_{HF}, and CE_W at the three basket points x^*_{HF}, x^*_W, x^*_R at $\omega = 0.50$. The crossing of CE_R and CE_{HF} is the Triffin dilemma in one image: moving from HF's basket to R's basket continuously improves R's welfare while destroying HF's welfare. The world welfare CE_W is maximized at the interior point x^*_W.

The right panel plots, for the two blocs R and HF, conditional welfare level curves under unilateral imposition—$CE_R(x^*_{HF};\omega)$ and $CE_{HF}(x^*_R;\omega)$ (solid)—and under coordination—$CE_R(x^*_W;\omega)$ and $CE_{HF}(x^*_W;\omega)$ (dashed)—as functions of adoption $\omega \in [0, 1]$. The shaded areas are the coordinator gains $\text{CG}_R(\omega)$ and $\text{CG}_{HF}(\omega)$—the welfare each bloc recovers when a benevolent planner imposes the world-welfare-optimal basket $x^*_W \approx (0.12, 0.34)$ rather than the opposing bloc's preferred basket. At $\omega = 0.50$, the coordinator recovers almost 90 percent of HF's conflict but less than 52 percent of R's—the world optimum leans safe-bloc, and the peripheral economy bears most of the irreducible adjustment cost (Table 4, Panel C).

The comparison with the original Triffin dilemma is precise. Triffin observed that the Bretton Woods system required the US to run gold-eroding deficits to supply the liquidity the world economy demanded. The two objectives—liquidity supply and gold convertibility—shared the same instrument and could not be simultaneously optimized. The new dilemma requires a global stablecoin to provide purchasing-power insurance to the periphery in a way that destabilizes the core through demand volatility; the two objectives again share the same instrument—the UoA basket. But the new dilemma has a dimension the original lacked: Proposition 4 guarantees that at the knife-edge the financial stability instrument (denomination) and the distributional instrument (UoA) are orthogonal. Denomination regulation can be imposed without touching the distributional question. This conditional

separability is what makes the new dilemma institutionally tractable, even if the political economy of resolution is not.

Scale intensification and forward-looking urgency. The right panel of Figure 8 shows that all welfare losses grow in magnitude as ω rises: the gap between each bloc's own-optimal welfare and its worst-case welfare widens monotonically with adoption. The coordination gains $\mathrm{CG}_j(\omega)$ also grow, but the residual losses grow faster for R. At the empirically current share of $\omega \approx 0.015$, welfare differences are quantitatively negligible. They become material around $\omega \approx 0.10$–0.20, a threshold reachable within a decade. Regulators who require observable welfare damage before acting will be regulating after the conflict has already entrenched—precisely the forward-looking failure the original Triffin dilemma illustrated.[7]

4.7 Robustness

The dominance of the unit-of-account margin documented in Table 2 is not an artifact of the benchmark calibration. Table 3 shows it is robust to low adoption and large parameter changes, and exhibits a hump-shaped amplification pattern across adoption levels.

Adoption. Panel A varies ω. Lower adoption compresses welfare levels: the absolute world-welfare range along the unit-of-account margin, ΔCE_W^U, declines by an order of magnitude as ω falls from 0.50 to 0.015. Yet the ranking of design margins is unchanged. The U/D ratio remains on the order of 10^2, exceeding 160× even at very low adoption. Adoption therefore scales the *level* of the design problem but not its *ranking*.

Nonmonotonic amplification. The strength of the hierarchy is nonmonotonic in adoption. The U/D ratio rises from 163× at $\omega = 0.015$ to a peak of 274× at $\omega = 0.35$, then falls back to 193× at the baseline $\omega = 0.50$. This hump-shaped pattern reflects curvature in how the two channels scale: at low and intermediate adoption the unit-of-account channel strengthens disproportionately, but beyond $\omega \approx 0.35$ the denomination channel also strengthens, compressing the ratio without overturning the ranking.

[7] This statement is conditional on exogenous adoption. The present framework characterizes transmission and welfare at a given adoption scale ω, but does not endogenize the take-up decision. In a richer model, the Triffin conflict itself could feed back into adoption and keep equilibrium ω below the threshold at which governance costs become worth incurring. Endogenizing adoption is therefore important for policy timing and for the quantitative magnitude of the conflict, though not for the conditional ranking of design margins at a given ω.

Preferences and financial amplification. Panel B varies the liquidity-gap weight φ_l, money-stablecoin substitution η, and the portfolio-balance coefficient ψ_b at fixed $\omega = 0.50$. Reducing the liquidity-gap weight to $\varphi_l = 0.10$ still yields a U/D ratio above 50×. Varying η and ψ_b similarly affects magnitudes but not the ordering. The world-welfare-maximizing basket weight $x_W^* = 0.35$ is invariant to all parameter changes considered, confirming it is a structural feature of the model rather than a calibration artifact.

Mechanism. The robustness follows from the model's equilibrium structure. The unit-of-account channel operates through the money-stablecoin demand wedge (11): deviations between stablecoin basket inflation π_j^{sc} and domestic CPI shift household demand between money and stablecoins, feeding directly into the IS curve. This channel reaches all three economies simultaneously and is active for any nonzero adoption ω. The denomination channel, by contrast, works indirectly through portfolio-balance premia and exchange-rate volatility, with strength proportional to the denomination imbalance δ_{jk}. At the knife-edge $\delta_{jk} = 0$ by construction, so the denomination channel is structurally suppressed precisely when the UoA channel operates at full strength. The hierarchy is therefore a direct consequence of the orthogonality established in Proposition 4.

5 Policy Implications

The findings in the previous sections imply a stablecoin regulatory framework built around two instruments, reflecting the orthogonality of the denomination and unit-of-account channels at the knife-edge benchmark and their interaction away from it. 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 target, implemented as a tolerance band around equal weights. Because financial spillovers scale linearly with denomination imbalances, maintaining θ^d within a neighborhood of $(1/3, 1/3, 1/3)$ is sufficient to substantially attenuate economically meaningful UIP distortions. The knife-edge therefore defines the center of a regulatory region rather than a point requirement. The requirement asks private issuers to reproduce by design the monetary neutrality that a CBDC achieves by the nature of sovereign issuance—making denomination diversification the private-sector analog of sovereign digital currency issuance from the perspective of international financial stability. 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 in the base-

line calibration. 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 harbor: 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 quantitatively similar 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 targets the world-welfare-optimal basket in the baseline calibration, $x_W^* \approx 0.34$, which Table 2 shows is stable across all architectures. This requires multilateral negotiation: the advanced-economy bloc prefers $x_{HF}^* \approx 0.09$ (near R-exclusion) and the periphery prefers $x_R^* \approx 0.87$ (heavy in own currency). The world optimum lies between these positions and cannot be implemented by any single party acting unilaterally. Unit-of-account governance—relating to pricing conventions and collateral frameworks—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 analog of SDR basket weights should look like in a world where private global liquidity provision has become quantitatively significant.

Denomination and unit of account are distinct instruments, but they are not distinct objectives: both operate on the same welfare margins through different transmission channels. Welfare non-separability implies that denomination regulation and unit-of-account governance cannot be optimized independently. The welfare-maximizing design reflects the joint determination of (x_d^*, x_u^*), so institutional arrangements that assign the two margins to separate authorities (e.g. FSB for denomination and IMF for the unit of account) must coordinate closely, as the marginal value of each instrument depends on the level of the other. This creates a case for either tightly coordinated governance across institutions or a single multilateral authority capable of jointly setting both dimensions.

Capital requirements complement but cannot replace denomination regulation. Figure 4 shows that the equity-feedback coefficient $\chi_E > 0$ is stabilizing under the knife-

edge but can amplify under concentrated denomination once the equity buffer is exhausted. A natural regulatory sequencing is: require knife-edge denomination first, then set capital requirements as a second line of defense 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 potentially more consequential for welfare, as the welfare analysis demonstrates. 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. At the knife-edge, denomination and unit-of-account governance enter disjoint equilibrium equations and can be separated into independent instruments. 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. This combination—denomination neutrality by construction, welfare-optimal basket by multilateral mandate—is precisely what the current global financial safety net lacks, and what the welfare non-separability result makes necessary. 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 welfare margin that is quantitatively smaller than the unit-of-account channel in the baseline calibration.

Stablecoin regulation that addresses denomination architecture without engaging with basket composition solves 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 at the knife-edge benchmark: 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 a strict hierarchy: the UoA channel generates welfare variation substantially larger in magnitude than the denomination channel along the risky-currency simplex edge (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 (Figure 8). The welfare-maximizing UoA for the periphery is substantially 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 unit-of-account governance are operationally separable instruments. At the knife-edge, they enter disjoint equilibrium equations and can therefore be delegated to different institutional mandates. The original Triffin dilemma lacked this property: liquidity provision and dollar credibility were tied to the same instrument.

The welfare objective, however, is not separable. The optimal design reflects a joint trade-off between insurance and distortion across the two margins, so the welfare-maximizing allocation requires their simultaneous choice. This creates a case for either tightly coordinated governance across institutions or a single multilateral framework—such as a digital SDR architecture—capable of jointly determining both denomination and unit of account.

Several extensions warrant future research. Endogenizing 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. 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 would resolve the New Triffin Dilemma by design 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 denomination 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 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 analog 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.

Nelson Camanho, Harald Hau, and Hélène Rey. Global portfolio rebalancing and exchange rates. *Review of Financial Studies*, 35(11):5228–5274, 2022.

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, 2015a.

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

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, 2021a.

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

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.

Oleg Itskhoki and Dmitry Mukhin. Exchange rate disconnect in general equilibrium. *Journal of Political Economy*, 129(8):2183–2232, 2021.

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

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

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.

Stephanie Schmitt-Grohé and Martín Uribe. Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, 28(4):755–775, 2004. doi: 10.1016/S0165-1889(03)00043-5.

Michael Woodford. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ, 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 $\dot{e}_{jk,t} \equiv e_{jk,t} - e_{jk,t-1}$. Superscript $^{+1}$ denotes one-period-ahead expectations, $\mathbb{E}_t x^{+1} \equiv \mathbb{E}_t x_{t+1}$. Country index $j \in \{H, F, R\}$; bilateral pair index $(jk) \in \{HF, HR, FR\}$ unless stated otherwise.

A.1 Supply.

$$\pi_{jd,t} = \beta \mathbb{E}_t \pi_{jd,t+1} + \kappa\, y_{j,t};$$
$$\kappa = \frac{(1-\alpha)(1-\beta\alpha)}{\alpha} \cdot \frac{1}{\mu}.$$

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.

$$\pi_{H,t} = (1-\gamma)\pi_{Hd,t} + \tfrac{\gamma}{2}(\pi_{Fd,t} + \dot{e}_{HF,t}) + \tfrac{\gamma}{2}(\pi_{Rd,t} + \dot{e}_{HR,t});$$
$$\pi_{F,t} = \tfrac{\gamma}{2}(\pi_{Hd,t} - \dot{e}_{HF,t}) + (1-\gamma)\pi_{Fd,t} + \tfrac{\gamma}{2}(\pi_{Rd,t} + \dot{e}_{FR,t});$$
$$\pi_{R,t} = \tfrac{\gamma}{2}(\pi_{Hd,t} - \dot{e}_{HR,t}) + \tfrac{\gamma}{2}(\pi_{Fd,t} - \dot{e}_{FR,t}) + (1-\gamma)\pi_{Rd,t}.$$

A.4 Terms of trade.

$$\dot{\text{tot}}_{HF,t} = \dot{e}_{HF,t} + \pi_{Fd,t} - \pi_{Hd,t};$$
$$\dot{\text{tot}}_{HR,t} = \dot{e}_{HR,t} + \pi_{Rd,t} - \pi_{Hd,t};$$
$$\dot{\text{tot}}_{FR,t} = \dot{e}_{FR,t} + \pi_{Rd,t} - \pi_{Fd,t};$$
$$\text{tot}_{FR,t} = \text{tot}_{HR,t} - \text{tot}_{HF,t}.$$

A.5–A.6 Liquidity services and liquidity gap.

$$\begin{aligned} l_{j,t} &= (1-\omega)\, mb_{j,t} + \omega\, s_{j,t}; \\ \text{gap}_{j,t} &= c_{j,t} - l_{j,t}. \end{aligned}$$

A.7 Euler equation.

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

A.8 Money demand.

$$mb_{j,t} = -\eta\, i_{j,t}.$$

A.9–A.11 Stablecoin basket inflation and money-stablecoin demand wedges.

$$\begin{aligned} \pi^{\text{sc}}_t &= \theta^u_H \pi_{Hd,t} + \theta^u_F(\pi_{Fd,t} + \dot{e}_{HF,t}) + \theta^u_R(\pi_{Rd,t} + \dot{e}_{HR,t}); \\ \pi^{\text{sc}}_{F,t} &= \pi^{\text{sc}}_t - \dot{e}_{HF,t}; \\ \pi^{\text{sc}}_{R,t} &= \pi^{\text{sc}}_t - \dot{e}_{HR,t}; \\ mb_{j,t} - s_{j,t} &= \eta(\pi^{\text{sc}}_{j,t} - \pi_{j,t}). \end{aligned}$$

A.12–A.13 Stablecoin aggregation and adoption.

$$\begin{aligned} s_t &= \lambda_s\, s_{H,t} + \tfrac{1-\lambda_s}{2}\, s_{F,t} + \tfrac{1-\lambda_s}{2}\, s_{R,t}; \\ \bar{s} &= \omega. \end{aligned}$$

A.14–A.16 UIP and exchange rates.

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

A.17–A.22 Risk premia.

$$
\begin{aligned}
\psi^{\text{sc}}_{HF,t} &= \psi_b\, \delta_{HF}\, \bar{s}(s_t - \zeta_E E^{\text{sc}}_t);\\
\psi^{\text{sc}}_{HR,t} &= \psi_b\, \delta_{HR}\, \bar{s}(s_t - \zeta_E E^{\text{sc}}_t);\\
\psi^{\text{sc}}_{FR,t} &= \psi^{\text{sc}}_{HR,t} - \psi^{\text{sc}}_{HF,t};\\
\psi_{HF,t} &= \psi_b\, b_{HF,t} + \psi^{\text{sc}}_{HF,t} - \chi_E E^{\text{sc}}_t;\\
\psi_{HR,t} &= \psi_b\, b_{HR,t} + \psi^{\text{sc}}_{HR,t} - \chi_E E^{\text{sc}}_t + \nu_t;\\
\psi_{FR,t} &= \psi_{HR,t} - \psi_{HF,t}.
\end{aligned}
$$

A.23 Risk-off shock.

$$
\nu_t = \rho_\nu\, \nu_{t-1} + \varepsilon^\nu_t, \qquad \varepsilon^\nu_t \sim \text{iid}(0, \sigma^2_\nu).
$$

A.24–A.29 Issuer block.

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

A.30–A.33 Net foreign assets.

$$
\begin{aligned}
b_{jk,t} &= \beta^{-1} b_{jk,t-1} + nx_{jk,t}, \quad (jk) \in \{HF,\, HR\};\\
b_{FR,t} &= b_{HR,t} - b_{HF,t};\\
nx_{jk,t} &= \tfrac{\gamma}{2}(c_{k,t} - c_{j,t}) + \tfrac{\gamma}{2}\,\text{tot}_{jk,t}, \quad (jk) \in \{HF,\, HR,\, FR\}.
\end{aligned}
$$

A.34–A.36 Goods markets and transfers.

$$
\begin{aligned}
T_{j,t} &= \omega_j\, \phi_{\text{div}} E^{\text{sc}}_{t-1}, \quad \omega_H = \lambda_s,\ \omega_F = \omega_R = \tfrac{1-\lambda_s}{2};\\
y_{j,t} &= c_{j,t} + nx_{j,t} - T_{j,t}.
\end{aligned}
$$

B 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 **fourteen** 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; two independent bilateral exchange rates $\{e_{HF}, e_{HR}\}$ from the UIP conditions (the F-R rate satisfies triangular no-arbitrage and is not a separate jump variable); three domestic inflation rates $\{\pi_{Hd}, \pi_{Fd}, \pi_{Rd}\}$ from the Phillips curves; 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. Issuer equity E^{sc} and bilateral NFA positions b_{jk} are predetermined state variables.

Dynare's `check` command confirms exactly **fourteen** 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^{sc}_{HF} to the ν shock are exactly zero by construction.

C Consumption-Equivalent Welfare

This appendix establishes that welfare can be evaluated from first-order dynamics and expressed as a function of stablecoin design (x_d, x_u). The central result is that separability is only an equilibrium property, not a welfare property: although denomination and unit-of-account enter disjoint equations at the knife-edge, the welfare objective is non-separable. The value of each margin depends on the other, so the optimum requires joint rather than sequential choice of both margins.

D.1 Second-Order Approximation of Utility

The representative household in country j maximizes:

$$U_j = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_{j,t}^{1-\sigma}}{1-\sigma} - \frac{\varphi_l}{2}(C_{j,t} - L_{j,t})^2 \right]. \tag{34}$$

Normalizing the steady state so that $\bar{C}_j = 1$, with $\bar{C} = \bar{L}$ at the symmetric steady state, a second-order Taylor expansion gives:

$$\frac{C_{j,t}^{1-\sigma}}{1-\sigma} \approx \frac{1}{1-\sigma} + c_{j,t} - \frac{\sigma}{2} c_{j,t}^2 + O(3). \tag{35}$$

For the liquidity wedge term, the level and linear terms vanish because $\bar{C} = \bar{L}$, leaving:

$$\frac{\varphi_l}{2}(C_{j,t} - L_{j,t})^2 \approx \frac{\varphi_l}{2} \operatorname{gap}_{j,t}^2 + O(3), \tag{36}$$

where $\operatorname{gap}_{j,t} = c_{j,t} - l_{j,t}$. Combining, period utility up to second order is:

$$u_{j,t} \approx \text{const} + c_{j,t} - \frac{\sigma}{2} c_{j,t}^2 - \frac{\varphi_l}{2} \operatorname{gap}_{j,t}^2 + O(3). \tag{37}$$

D.2 Price Dispersion and Productive Inefficiency

Under Calvo pricing, the cross-sectional dispersion of prices creates productive inefficiency. Following Woodford (2003), Chapter 6, the second-order approximation of the price dispersion index gives:

$$\log \Delta_{j,t} \approx \frac{\mu}{2(1-\alpha)(1-\beta\alpha)} \pi_{jd,t}^2 = \frac{1}{2\kappa} \pi_{jd,t}^2. \tag{38}$$

Substituting the dispersion cost into the utility expansion yields the additional welfare-loss term $-\frac{\kappa}{2\mu}\pi_{jd,t}^2$, with welfare weight $\kappa/\mu = (1-\alpha)(1-\beta\alpha)/(\alpha\mu^2)$. The inflation stabilization motive therefore originates on the supply side through productive inefficiency.

D.3 Eliminating Linear Terms

Under the symmetric steady state and the assumed policy rules, the discounted sum of linear deviations satisfies:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t c_{j,t} = \text{t.i.p.} + O(2), \tag{39}$$

where t.i.p. collects terms independent of the stablecoin design parameters. The welfare loss relative to the efficient allocation is therefore:

$$\mathcal{W}_j \approx -\frac{1}{2} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\sigma\, c_{j,t}^2 + \varphi_l \operatorname{gap}_{j,t}^2 + \frac{\kappa}{\mu}\, \pi_{jd,t}^2 \right] + \text{t.i.p.} \tag{40}$$

The inflation weight $\kappa/\mu = 0.021$ is small relative to σ and φ_l at the baseline, confirming that the inflation channel is quantitatively secondary.

D.4 Consumption-Equivalent Metric

The consumption-equivalent welfare difference between Regime A and Regime B for country j is:

$$\Delta CE_j \approx \frac{\mathcal{W}_j^A - \mathcal{W}_j^B}{1/(1-\beta)}, \tag{41}$$

which scales the discounted welfare difference by the steady-state utility weight. World and safe-bloc CE welfare are aggregated as:

$$CE_{\text{world}} = \tfrac{1}{3}(CE_H + CE_F + CE_R), \qquad CE_{HF} = \tfrac{1}{2}(CE_H + CE_F). \tag{42}$$

D.5 Evaluation from First-Order Dynamics

The second-order welfare object in (40) can be evaluated exactly from the first-order solution (Woodford, 2003; Schmitt-Grohé and Uribe, 2004). For any variable $x_{j,t}$ linear in the state vector:

$$M_x^2 \equiv \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t x_{j,t}^2 = \sum_{h=0}^{\infty} \beta^h \left[\text{IRF}_x(h)\right]^2, \tag{43}$$

where $\text{IRF}_x(h)$ is the h-period impulse response computable directly from the Dynare first-order solution matrices. The CE welfare difference is:

$$\Delta CE_j = -\frac{1-\beta}{2}\left[\sigma\left(M_c^{2,A} - M_c^{2,B}\right) + \varphi_l\left(M_{\text{gap}}^{2,A} - M_{\text{gap}}^{2,B}\right) + \frac{\kappa}{\mu}\left(M_\pi^{2,A} - M_\pi^{2,B}\right)\right]. \tag{44}$$

A positive ΔCE_j means Regime A delivers higher welfare than Regime B.

D.6 Validity Conditions

The linear-quadratic approach is valid under three conditions, all satisfied here: zero-mean shocks and symmetric steady state; policy rules independent of second-order terms; and design parameters entering exclusively through first-order dynamics. Under these conditions, equation (44) provides an exact welfare ranking at second order from first-order dynamics.

D.7 Quantitative Welfare Analysis

Table 4 summarizes the quantitative welfare results underpinning Section 4.6. Panel A reports optimal designs, mapping onto the white star and bloc optima of Figure 7. Panel B decomposes welfare loss, the numerical counterpart to Figure 7. Panel C quantifies conflict and coordination at $\omega = 0.50$, the exact snapshot of the right panel of Figure 8.

Optimal designs (Panel A). The world optimum at $(x_d^*, x_u^*) = (0.12, 0.34)$ is interior in both dimensions. The HF optimum at $(0.00, 0.09)$ excludes R entirely from denomination while allowing a small R weight in the unit of account. The R optimum at $(0.19, 0.87)$ is the most revealing: R wants a strongly R-inclusive UoA but a substantially lower R-denomination weight ($x_d^* = 0.19$, far from the corner $x_d^* = 1.00$). The gap between R's denomination preference at knife-edge UoA ($x_d^* = 1.00$ at $\bar{x}_u = 0.33$) and at its own unit-of-account optimum ($x_d^* = 0.19$ at $x_u^* = 0.87$) is the first signal of welfare non-separability: the higher unit-of-account margin already provides the purchasing-power insurance that denomination would otherwise supply.

Welfare-loss decomposition (Panel B). Three results stand out. First, the liquidity-gap variance M_{gap}^2 is the largest raw second moment at the world optimum (578×10^{-3} versus 156×10^{-3} for consumption), confirming that the stablecoin functions primarily as purchasing-power insurance through the money-demand channel. Second, once utility weights are applied, consumption volatility dominates world welfare loss ($s_c = 61.8\%$): the weight $\sigma = 3$ amplifies consumption variance relative to $\varphi_l = 0.5$ on the gap. Third, inflation contributes less than 0.1% at every optimum: the New Keynesian price-dispersion channel plays no role in the welfare ranking across stablecoin architectures. Optimal stablecoin design is thus the solution to an insurance-distortion trade-off. Denomination provides insurance against exchange-rate-induced fluctuations in domestic purchasing power, but only through a distorted financial channel—the UIP wedge—which feeds into exchange rates, the output gap, and ultimately consumption volatility. The unit-of-account margin delivers insurance without this distortion. HF therefore excludes R from the unit of account ($x_u^* = 0.00$), while R relies primarily on the UoA margin ($x_u^* = 0.87$) and limits R-denomination ($x_d^* = 0.19$), balancing insurance against distortion. The two margins are thus not independently optimizable: the optimal design is jointly determined by the insurance-distortion trade-off across denomination and unit of account.

Conflict and coordinator gain (Panel C). Panel C quantifies the Triffin conflict at $\omega = 0.50$. Being forced to the opposing bloc's preferred basket costs R $\text{Loss}_R = 6.67 \times 10^{-3}$—2.3 times R's own-optimum welfare level—and costs HF $\text{Loss}_{HF} = 7.23 \times 10^{-3}$—7.5 times its own-optimum level. A benevolent coordinator imposing the world-welfare-optimal basket $x_W^* = (0.12, 0.34)$ recovers $\text{CG}_R = 3.45 \times 10^{-3}$ for R (52 percent of its total loss) and $\text{CG}_{HF} = 6.50 \times 10^{-3}$ for HF (90 percent of its total loss). The residual loss after coordination—48 percent of R's total loss and 10 percent of HF's total loss—is the conflict the coordinator cannot resolve. The world optimum leans toward the safe bloc, and R bears most of the

irreducible adjustment cost. This asymmetry is the quantitative essence of the new Triffin dilemma: even under optimal multilateral governance, the peripheral economy absorbs the majority of the distributional conflict.

Welfare non-separability. Table 5 demonstrates quantitatively the welfare non-separability between the two design dimensions: it is the numerical counterpart to the tilted contour geometry in Figure 7.

Three results follow directly from comparing the two panels. First, the unit-of-account margin dominates conditionally: at the knife-edge, $\Delta CE_R^U = 3.23 \times 10^{-3}$ versus $\Delta CE_R^D = 0.37 \times 10^{-3}$, a ratio of 8.8×. Second, this dominance is robust across denomination choices: ΔCE_R^U varies by only 20 percent over the full denomination range (Panel A), whereas ΔCE_R^D varies by 98 percent across the unit-of-account range (Panel B). Third, the two instruments are welfare substitutes for R: as $\bar{x}_u$ approaches R's preferred unit of account, the marginal welfare value of denomination collapses—ΔCE_R^D falls from 0.369×10^{-3} to 0.007×10^{-3} as $\bar{x}_u$ rises from 1/3 to 0.87. Strikingly, when the unit of account excludes R entirely ($\bar{x}_u = 0$), R's conditional optimum over denomination lies at the H corner rather than the R corner. This is because, at $\bar{x}_u = 0$, full R-denomination would generate substantial macroeconomic volatility through the UIP wedge. Although denomination provides insurance against R depreciation, this insurance operates through a distorted financial channel that amplifies fluctuations in consumption and the output gap. In equilibrium, the distortion dominates the benefit, so even R households do not choose R-denominated stablecoins on that slice—reflecting the underlying insurance-distortion trade-off in the model.

Together, the panels confirm welfare non-separability and establish that UoA is the primary welfare instrument while denomination, though essential for financial stability, is the secondary welfare lever.

D Proofs

Proof of Proposition 1. Substituting $\theta_H^d = \theta_F^d = \theta_R^d = 1/3$ into (A.17)–(A.19): each concentration coefficient $(1/3 - 1/3) = 0$, so $\psi_{HF,t}^{\text{sc}} = \psi_{HR,t}^{\text{sc}} = \psi_{FR,t}^{\text{sc}} = 0$ for all t, all s_t, E_t^{sc}, and all shock realizations. The result holds for any ψ_b, $\bar{s}$, ζ_E, ω. □ □

Remark 1 (Structural basis of denomination neutrality)**.** *The exact-zero outcome in Proposition 1 is not a calibration artifact but a structural property of the portfolio-balance microfoundation. Three features are essential. First, the currency exposure of the stablecoin is determined by the denomination vector θ^d on the liability side of the issuer's balance sheet, independently of the distribution of holdings across countries. Second, the uniform*

ψ_b follows from the symmetry of the variance-covariance matrix under the symmetric three-country calibration; relaxing this assumption would deliver bilateral-specific coefficients $\psi_{b,jk}$ but would not restore the contagion channel shut down by the knife-edge, since the net imbalance $(\theta_j^d - \theta_k^d)$ would still be zero. Third, the linearity of the premium in $(\theta_j^d - \theta_k^d)$ follows from the linearity of mean-variance portfolio demand in net positions: no approximation is involved at first order.

Proof of Corollary 1. The "if" direction follows immediately from Proposition 1: under the knife-edge, all denomination-channel UIP wedges are identically zero. The "only if" direction follows from equations (A.17)–(A.19): the wedge $\psi^{\text{sc}}_{HF,t} = \psi_b(\theta_H^d - \theta_F^d)\bar{s}(s_t - \zeta_E E_t^{\text{sc}})$ is zero for all s_t if and only if $\theta_H^d = \theta_F^d$, and similarly for the other two bilaterals, which requires $\theta_H^d = \theta_F^d = \theta_R^d = 1/3$. □ □

Remark 2 (CBDC neutrality versus knife-edge neutrality)**.** *A CBDC issued by country j achieves monetary neutrality by a different mechanism. As a direct liability of the sovereign monetary authority, it is part of the monetary base and does not interpose a private intermediary between domestic money demand and international asset markets. The portfolio-balance wedge formula (A.17)–(A.19) therefore does not apply to a CBDC: neutrality holds by institutional definition rather than by parameter choice. The knife-edge stablecoin achieves the same neutrality outcome through a different route: the wedge formula applies because it is a private instrument, but the denomination vector is chosen such that $(\theta_j^d - \theta_k^d) = 0$ for all j, k, zeroing the wedge by design.*

Proof of Proposition 2. Under concentrated denomination $\theta_H^d = 1$, $\theta_F^d = \theta_R^d = 0$: $\psi^{\text{sc}}_{HF,t} = \psi_b \cdot \bar{s} \cdot s_t$. The risk-off shock ν_t enters $\psi_{HR,t}$ directly via (A.22), inducing $\Delta s_t \neq 0$ through the money-demand system; this change enters $\psi^{\text{sc}}_{HF,t}$ with positive weight and shifts $e_{HF,t}$ through (A.14). Under $\theta_H^d = \theta_F^d = \theta_R^d = 1/3$, $\psi^{\text{sc}}_{HF,t} = 0$ identically by Proposition 1, severing the transmission channel entirely. □ □

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, 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 in the H-per-R convention): $(\theta_R^a - \theta_R^d)\dot{e}_{HR,t} < 0$, giving $\text{val}_{fx,t} < 0$: equity is destroyed procyclically. □ □

Proof of Proposition 4. Part (i). Under the knife-edge, $\theta_H^d = \theta_F^d = \theta_R^d = 1/3$, so $\psi^{\text{sc}}_{jk} = \psi_b(\theta_j^d - \theta_k^d)\bar{s}(\cdot) = 0$ identically by Proposition 1. The coefficient on s_t in every UIP wedge is zero, so no path from θ^u through π^{sc} and s_t can affect any premium or exchange rate at any order. Exact orthogonality follows.

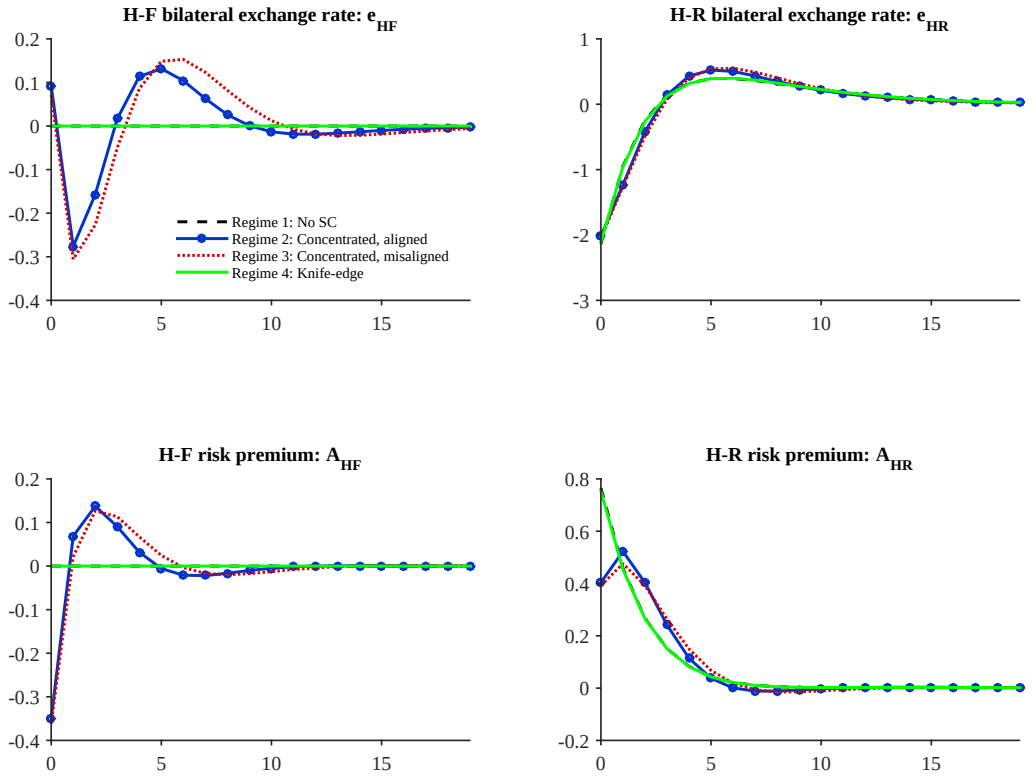
Part (ii). For general θ^d, θ^u enters the system exclusively through π^{sc} in (A.9). Tracing downstream: π^{sc} shifts $(mb_j - s_j)$ via (A.11), which shifts individual holdings s_j and hence aggregate s_t. Aggregate s_t enters the UIP wedges in (A.17)–(A.18) with slope $\psi_b(\theta^d_j - \theta^d_k)\bar{s}$, which does not depend directly on θ^u. True equilibrium independence of ψ^{sc}_{jk} with respect to θ^u therefore holds if and only if $\theta^d_j = \theta^d_k$ for all j, k. Off the knife-edge, θ^u affects ψ^{sc}_{jk} through s_t; the magnitude of this indirect effect is bounded by $\psi_b|\theta^d_j - \theta^d_k|\bar{s} \cdot |\partial s_t/\partial\theta^u|$ and represents a quantitatively negligible leakage from the unit-of-account channel into the financial block relative to its primary transmission through money demand, as confirmed by Table 2. □ □

Proof of Proposition 5. Let $CE_j(x_d, x_u)$ be the welfare objective defined in equation (44). It is sufficient to show that at least one of the equilibrium objects $(c_{j,t}, \text{gap}_{j,t}, \pi_{jd,t})$ has a non-zero mixed derivative with respect to (x_d, x_u). Along the H–R edge, define $\theta^d(x_d) = (1 - x_d,\, 0,\, x_d)$ and $\theta^u(x_u) = (1 - x_u,\, 0,\, x_u)$. The unit-of-account enters through stablecoin inflation: $\partial\pi^{\text{sc}}_t/\partial x_u = (\pi_{Rd,t} + \dot{e}_{HR,t}) - \pi_{Hd,t}$. This term enters the money-demand block through (A.11), implying $\partial s_t/\partial x_u \neq 0$ generically. Along the H–R edge, $\psi^{\text{sc}}_{HR,t} = \psi_b(1 - 2x_d)\bar{s}(s_t - \zeta_E E^{\text{sc}}_t)$, so:

$$\frac{\partial^2 \psi^{\text{sc}}_{HR,t}}{\partial x_d\, \partial x_u} = -2\psi_b\bar{s}\left(\frac{\partial s_t}{\partial x_u} - \zeta_E \frac{\partial E^{\text{sc}}_t}{\partial x_u}\right) \neq 0$$

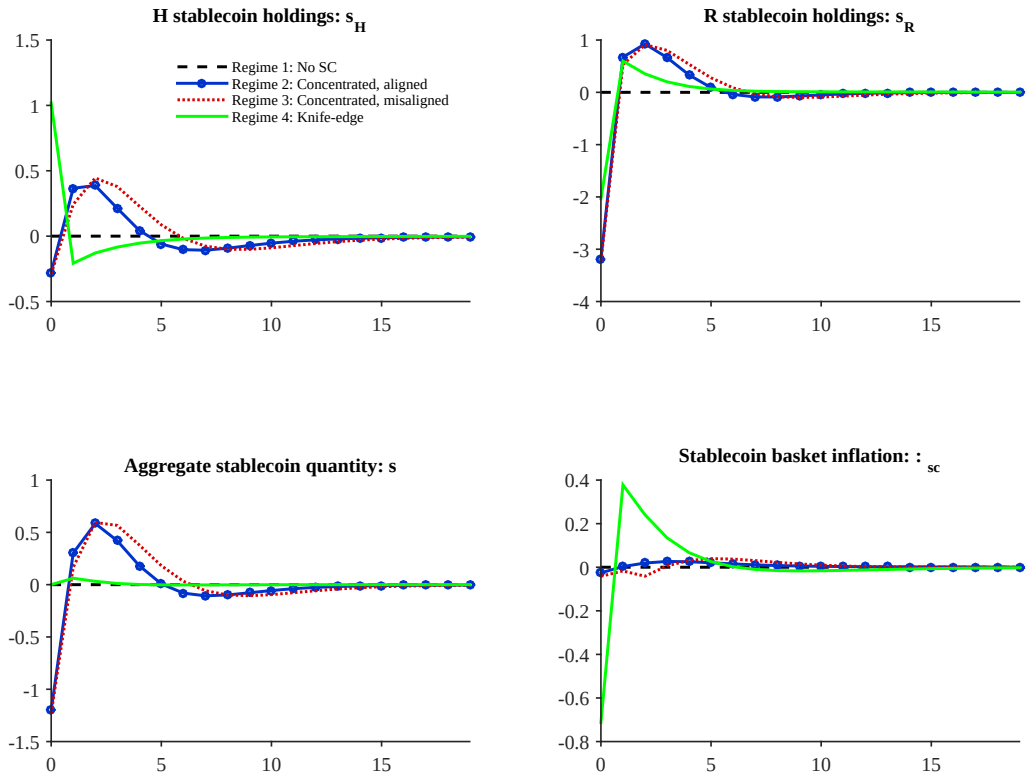
generically off the knife-edge. Since the UIP wedge enters exchange rates through (A.14), and exchange rates enter the IS and goods-market blocks, the mixed derivative of $c_{j,t}$ and $\text{gap}_{j,t}$ with respect to (x_d, x_u) is non-zero. Hence $\partial^2 CE_j/\partial x_d\, \partial x_u \neq 0$ generically, and $CE_j(x_d, x_u)$ cannot be written as $f(x_d) + g(x_u)$. Sequential optimization does not recover the joint optimum. □ □

Figure 1: Denomination Neutrality and Contagion



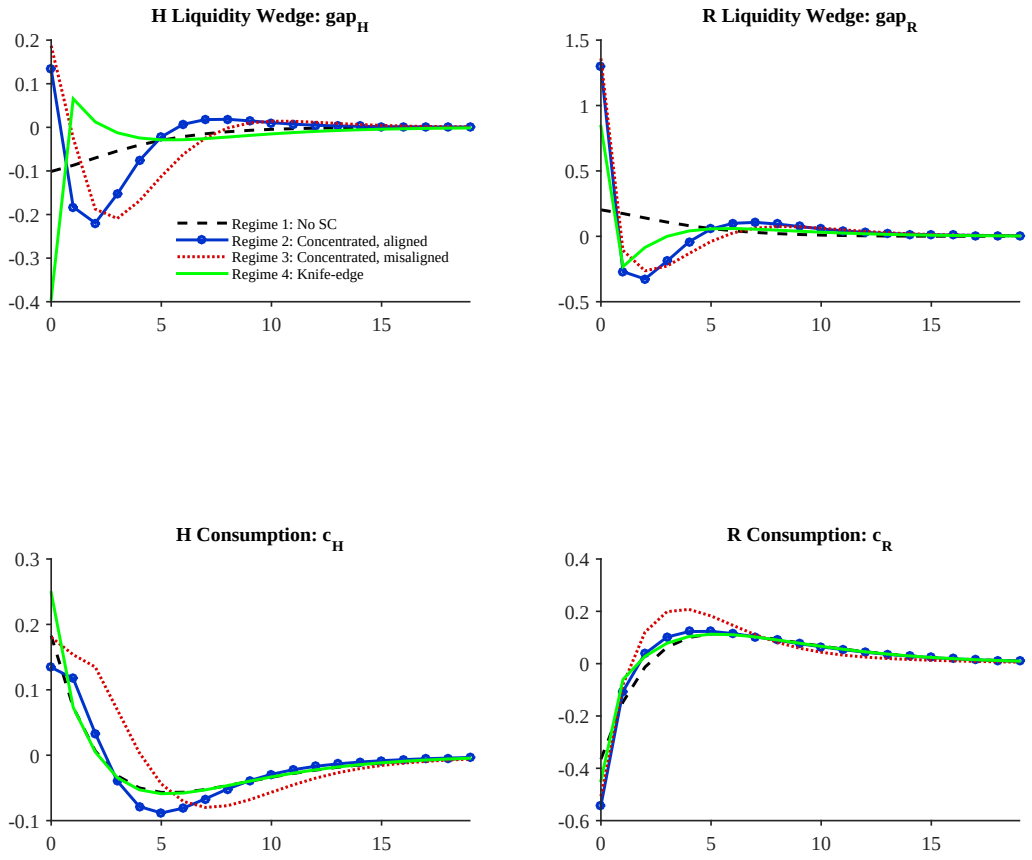
Note: Impulse responses to a one-standard-deviation global risk-off shock ν, entering ψ_{HR} only; the F-R premium inherits the shock residually through triangular no-arbitrage. Top: H-F bilateral exchange rate e_{HF} (left) and H-R rate e_{HR} (right). Bottom: H-F risk premium ψ_{HF} (left) and H-R 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.

Figure 2: The Transmission Chain



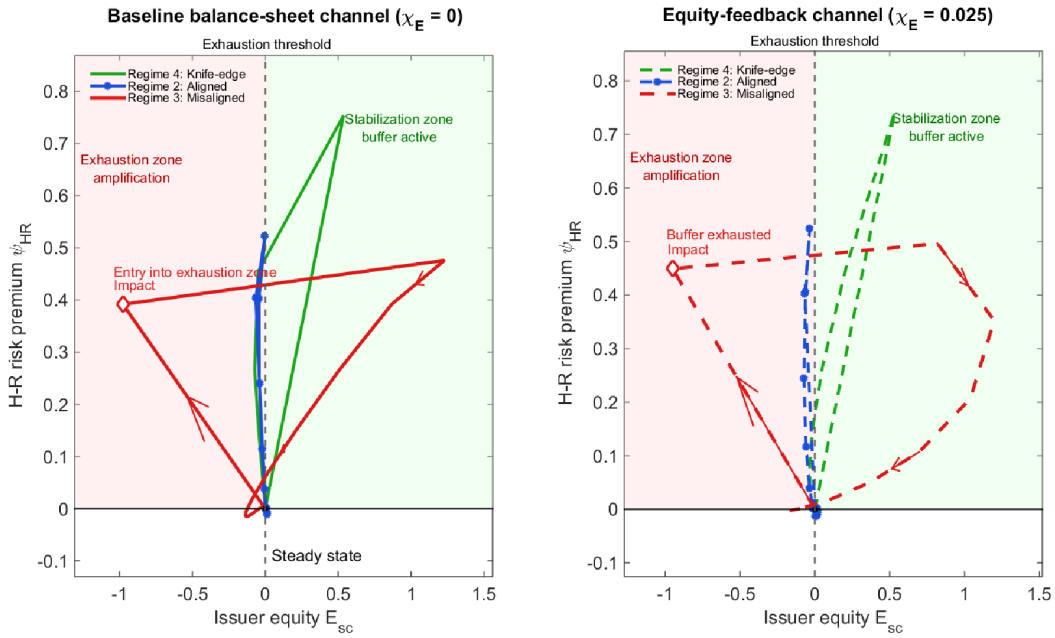
Note: Impulse responses to the same risk-off shock as Figure 1. Top-left: H-country stablecoin holdings s_H. Top-right: R-country holdings s_R. Bottom-left: aggregate stablecoin quantity s. Bottom-right: basket inflation π^{sc}. On impact, R households generate a brief safe-haven spike in s_R; by period 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 $\psi^{\text{sc}}_{HF} = \psi_b \, \delta_{HF} \, \bar{s} \cdot s$ (with $\delta_{HF} = 1$ under $\theta^d = (1, 0, 0)$ and $\zeta_E = 0$ at baseline) with a unit denomination-imbalance coefficient, delivering the premium spike of Figure 1. The basket inflation panel π^{sc} (bottom-right) illustrates Proposition 4: 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. 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.

Figure 3: The Real Consequences



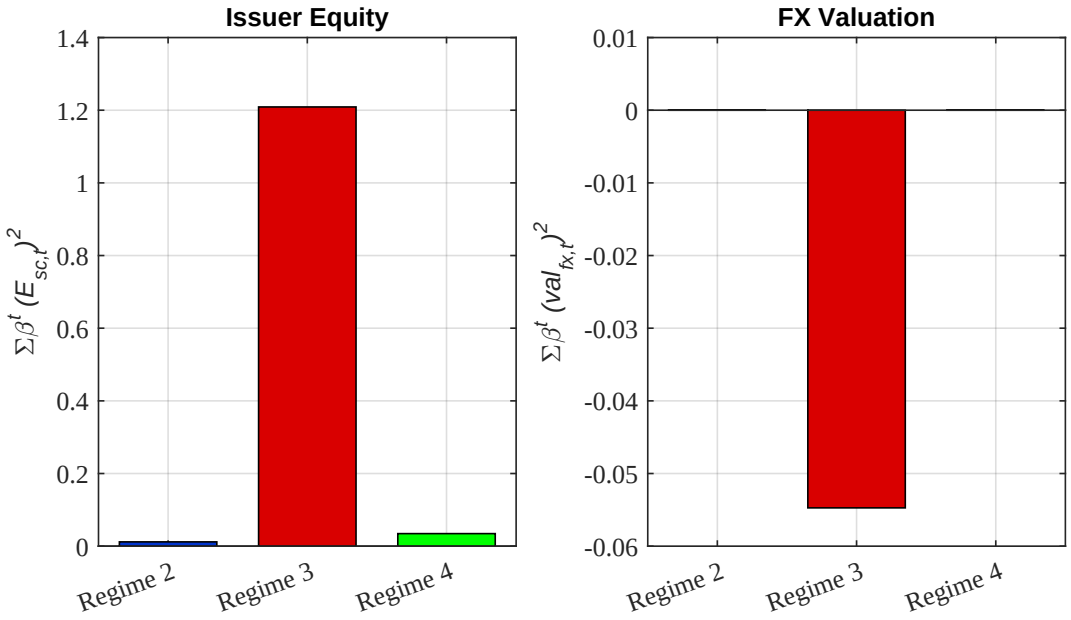
Note: Impulse responses to the same risk-off shock as Figure 1. Top row: H liquidity wedge $\text{gap}_H = c_H - l_H$ (left) and R liquidity wedge gap_R (right). Middle row: H consumption c_H (left) and R consumption c_R (right). Bottom row: H net exports nx_H (left) and R net exports nx_R (right). The top row shows the denomination channel hitting gap_H directly (Regime 4 is positive on impact, Regimes 2/3 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 Regimes 2/3 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.

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



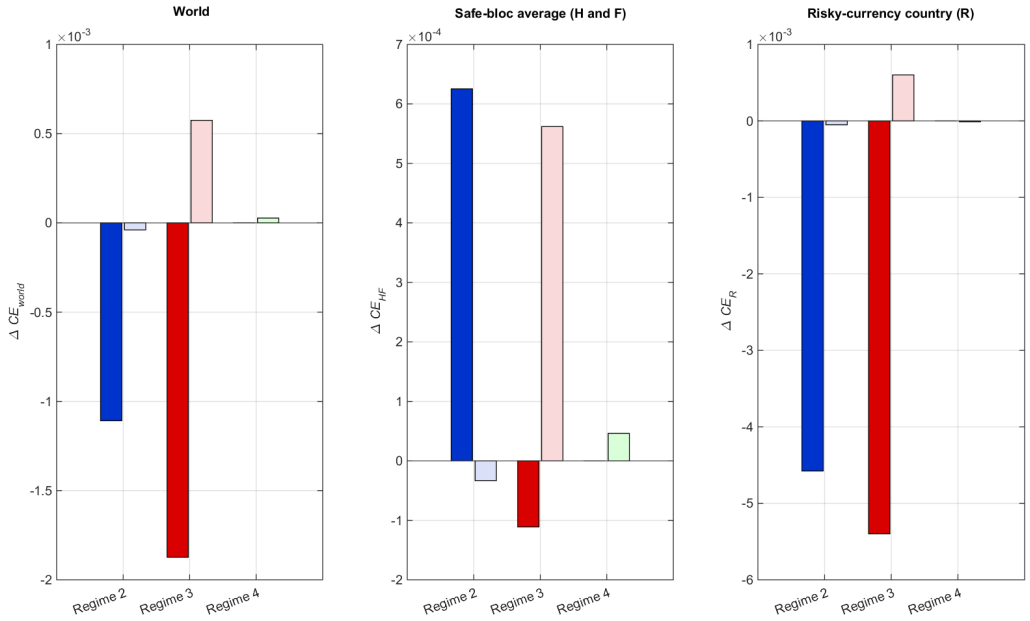
Note: Phase diagram space: issuer equity E^{sc} (x-axis) versus H-R risk premium ψ_{HR} (y-axis). Equity–premium dynamics under alternative architectures. Each panel plots impulse-response trajectories in the (E^{sc}, ψ_{HR}) space. The dashed vertical line at $E^{sc} = 0$ is a reference locus, not a constraint: to its right, equity dampens premia; to its left, it amplifies them through the linear feedback term $-\chi_E E^{sc}$. The figure visualizes how sign-dependent amplification arises within the linearized system. Trajectories crossing $E^{sc} = 0$ reflect continuous state evolution rather than a discrete regime shift or exhaustion threshold. 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 alignment eliminates the val_{fx} channel entirely. Regime 2 (blue) makes a smaller stabilization-zone loop. Regime 3 (red) enters the exhaustion zone on impact driven purely by val_{fx} destruction; the entire orbit lies in the red zone with no feedback required. *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: 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.

Figure 5: Buffer Effectiveness: The Magnitude of the Balance-Sheet Mechanism



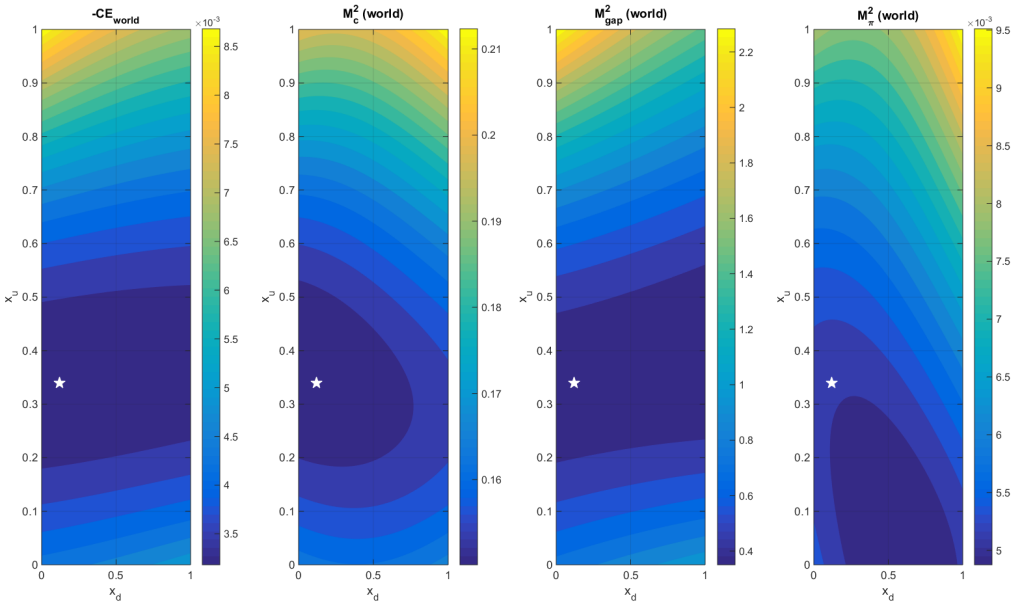
Note: Left panel: change in issuer equity discounted second moment $\sum_t \beta^t (E_t^{\text{sc}})^2$ from activating buffers, by regime. Right panel: change in FX valuation discounted second moment $\sum_t \beta^t (\text{val}_{fx,t})^2$ from activating buffers. Under Regime 3, buffers sharply increase equity volatility (left bar $\approx +1.2$) while reducing val_{fx} volatility (right bar ≈ -0.055). 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.

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



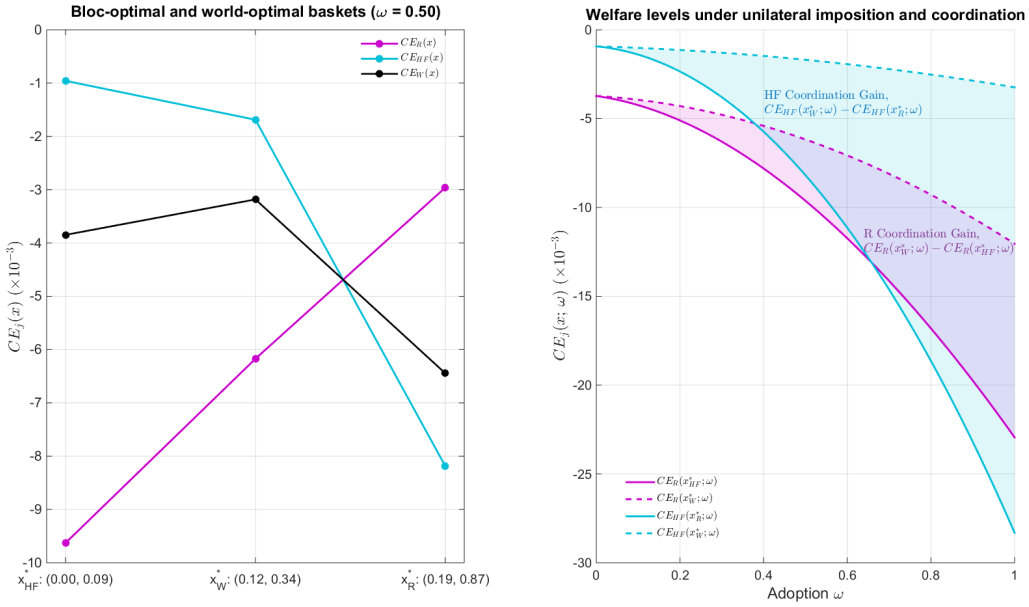
Note: Three country perspectives. Solid bars: architecture cost (CE welfare of Regime 2 or 3 relative to Regime 4, buffers off). Pale bars: incremental buffer effect (activating $\phi_{\text{div}} = 0.5$, $\chi_E = 0.025$, $\zeta_E = 0.05$ within each regime). Left: world welfare ΔCE_W ($\times 10^{-3}$). Center: safe-bloc average ΔCE_{HF} ($\times 10^{-4}$—note different scale). Right: R-country welfare ΔCE_R ($\times 10^{-3}$). The center panel shows a 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 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. The architecture most beneficial to the safe bloc is the same one most harmful to the periphery.

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



Note: Each panel varies denomination (x_d, R-currency weight on the x-axis) and unit of account (x_u, R-currency weight on the y-axis) *jointly* over $[0,1]^2$, with backing aligned to denomination throughout. The white star marks the two-dimensional welfare-maximizing design $(x_d^*, x_u^*) \approx (0.12, 0.34)$ in all panels. *Panel 1 (left):* world welfare loss. Contours are nearly horizontal, confirming the supremacy of the unit-of-account dimension. *Panel 2:* consumption conditional second moment. Contours display pronounced interior curvature, revealing denomination's distinct contribution to consumption volatility. *Panel 3:* liquidity-gap conditional second moment. The contour geometry closely mirrors Panel 1. *Panel 4 (right):* inflation conditional second moment ($\times 10^{-3}$). The surface is smooth and its range is small relative to Panels 1–3.

Figure 8: The New Triffin Dilemma



Note: conditional welfare and distributional conflict at the full two-dimensional optimal baskets. *Left panel:* static cross-section at $\omega = 0.50$. The horizontal axis indexes the three optimal baskets: the HF-optimal basket x^*_{HF}, the world-optimal basket x^*_W, and the R-optimal basket x^*_R. Each basket is defined by its denomination and unit-of-account weights on the R currency, (x_d, x_u), mapped into portfolio shares as $\theta_d(x_d) = \big((1-x_d)/2, (1-x_d)/2, x_d\big)$, $\theta_u(x_u) = \big((1-x_u)/2, (1-x_u)/2, x_u\big)$ with backing aligned to denomination, $\theta_a = \theta_d$. The three points correspond to the full two-dimensional optima $(x_d, x_u) \in \{x^*_{HF}, x^*_W, x^*_R\}$ obtained from the contour search (Table 4, Panel A). Each marker reports conditional welfare $CE_j(x;\omega)$ for all blocs evaluated at the same basket. The ordering across the three baskets reflects the underlying distributional conflict across blocs, with each bloc attaining its maximum at its own preferred basket and no single basket jointly maximizing welfare, giving rise to the Triffin trade-off. *Right panel:* CE welfare levels as functions of adoption ω, evaluated at the three two-dimensional basket points from Table 4, Panel A. Solid magenta: $CE_R(x^*_{HF};\omega)$—R's welfare under the worst imposition (forced to HF basket). Dashed magenta: $CE_R(x^*_W;\omega)$—R's welfare under the coordinator. Solid cyan: $CE_{HF}(x^*_R;\omega)$—HF's welfare under the worst imposition. Dashed cyan: $CE_{HF}(x^*_W;\omega)$—HF's welfare under the coordinator. Shaded areas are the coordinator gains $\mathrm{CG}_j(\omega)$ (equations (32)–(33)), the welfare each bloc recovers by moving from the worst unilateral imposition to the world-welfare-optimal basket, $x^*_W \approx (0.12, 0.34)$. The numerical snapshot at $\omega = 0.50$ is reported in Panel C of Table 4.

Table 1: Baseline Calibration

Parameter	Description	Value	Source / Target
Preferences and technology			
β	Discount factor	0.99	$r \approx 4\%$ ann. (AM calibration)
σ	Risk aversion	3.0	DSGE standard
η	Interest semi-elasticity	2.0	Money demand
φ_l	Liquidity wedge strength	0.50	SC adoption impact
Firms and monetary policy			
κ	NKPC slope	0.125	Calvo $\alpha = 0.75$, $\mu = 6$
α	Calvo price stickiness	0.75	4-quarter price duration
μ	Substitution elasticity	6	20% steady-state markup
κ/μ	Inflation welfare weight	0.021	Implied by κ, μ
γ	Import share	0.30	Trade data
ρ_i	Taylor smoothing	0.80	Standard
ϕ_π	Taylor inflation	1.50	Taylor principle
ϕ_y	Taylor output	1.00	Standard
International asset markets			
ψ_b	Portfolio-balance coeff.	0.60	Standard three-country value
ρ_ν	Risk premium persistence	0.80	VIX persistence
Stablecoin design (baseline)			
ω	SC adoption share	0.50	Scenario
$\bar{s}$	Steady-state SC supply	0.50	$\bar{s} = \kappa_S \omega$, $\kappa_S = 1$
κ_S	SC supply normalisation	1.00	Normalisation
λ_s	H share of SC holdings	1/3	Symmetric
ϕ_{div}	Issuer payout rate	1.00	Full pass-through baseline
χ_E	Equity-to-premium coeff.	0.00	Baseline (buffers off)
ζ_E	Equity-to-exposure coeff.	0.00	Baseline (buffers off)
Buffer regime (Section 4.5 only)			
ϕ_{div}	Payout rate (buffers on)	0.50	Partial retention
χ_E	Equity feedback (buffers on)	0.025	Calibrated to doom-loop
ζ_E	Equity exposure (buffers on)	0.05	Calibrated to doom-loop

Table 2: Welfare Architecture: Importance and Distributional Conflict

Regime	Margin	Welfare Range ($\times 10^3$)				Argmax: HR edge		
		HF edge		HR edge		World	Safe bloc	R
		ΔCE_W	ΔCE_R	ΔCE_W	ΔCE_R	x^*_W	x^*_{HF}	x^*_R
R2: Concentrated, aligned	D	0.52	0.88	0.91	2.26	0.00	0.12	0.00
	A	0.05	0.03	0.78	0.83	0.00	0.00	0.00
	U	**0.49**	**0.86**	**6.02**	**7.75**	**0.33**	**0.09**	**0.89**
R3: Concentrated, misaligned	D	0.20	0.36	0.20	0.73	0.53	0.88	0.30
	A	0.05	0.03	0.78	0.83	0.00	0.00	0.00
	U	**0.11**	**0.36**	**6.03**	**8.64**	**0.37**	**0.13**	**0.89**
R4: Knife-edge	D	*0.00*	*0.00*	0.03	0.34	0.69	0.00	1.00
	A	*0.00*	*0.00*	0.29	0.16	0.26	0.17	0.45
	U	***0.00***	***0.00***	**4.76**	**8.43**	**0.35**	**0.09**	**0.87**

Notes. ΔCE_W and ΔCE_R: range of world and R-country CE welfare (max−min) as the design parameter traverses the simplex edge ($\times 10^3$); world welfare $CE_W = \frac{1}{3}(CE_H + CE_F + CE_R)$ and safe-bloc welfare $CE_{HF} = \frac{1}{2}(CE_H + CE_F)$ per equation (42). Argmax columns report the HR-edge value $x^* \in [0, 1]$ that maximizes CE welfare for each country group, where $\theta(x) = (1 - x,\, 0,\, x)$, so that $x = 0$ is the full H-currency corner and $x = 1$ is the full R-currency corner. All HR-edge argmax values are computed holding the complementary margin at the knife-edge value 1/3. **Bold** rows: unit-of-account (U) margin. *Italic* entries: range $< 0.01 \times 10^{-3}$, economically negligible. U/D welfare ratios (HR edge, world welfare): R2 = 6.6×, R3 = 29.7×, R4 = 193×.

Table 3: Sensitivity of the Welfare Hierarchy to Adoption and Structural Parameters

		Peripheral CE welfare ($\times 10^3$)			HR edge summary		
Panel	**Value**	CE_R^{R2}	CE_R^{R3}	CE_R^{R4}	ΔCE_W^U	U/D ratio	x_W^*
Panel A. Adoption sensitivity							
ω	0.015	−3.79	−3.79	−3.75	0.06	163×	0.35
ω	0.05	−3.98	−3.97	−3.80	0.08	167×	0.35
ω	0.10	−4.35	−4.35	−3.92	0.11	175×	0.35
ω	0.20	−5.41	−5.46	−4.26	0.19	205×	0.35
ω	0.35	−7.70	−8.01	−5.05	0.31	274×	0.35
ω	0.50	−10.70	−11.60	−6.16	0.48	193×	0.35
Panel B. Parameter sensitivity at $\omega = 0.50$							
φ_l	0.10	−5.36	−5.69	−3.78	0.11	50.8×	0.35
φ_l	0.25	−7.65	−8.19	−4.78	0.26	113×	0.35
φ_l	0.50	−10.70	−11.60	−6.16	0.48	193×	0.35
φ_l	1.00	−15.10	−16.30	−8.13	0.89	299×	0.35
φ_l	2.00	−19.90	−21.60	−10.40	1.50	413×	0.35
η	1.00	−5.27	−5.44	−3.65	0.19	77.3×	0.35
η	1.50	−7.65	−8.11	−4.71	0.42	192×	0.35
η	2.00	−10.70	−11.60	−6.16	0.48	193×	0.35
η	3.00	−18.40	−20.30	−10.10	0.39	108×	0.35
η	4.00	−27.50	−30.70	−15.10	0.31	72.1×	0.35
ψ_b	0.10	−48.00	−48.90	−28.70	1.00	67.5×	0.35
ψ_b	0.30	−21.00	−22.20	−12.10	0.82	135×	0.35
ψ_b	0.60	−10.70	−11.60	−6.16	0.48	193×	0.35
ψ_b	0.90	−6.78	−7.35	−4.01	0.23	103×	0.35
ψ_b	1.20	−4.73	−5.13	−2.92	0.14	74.1×	0.35

Notes. CE_R^{Rk} denotes peripheral-country CE welfare under Regime $k \in \{2, 3, 4\}$, in $\times 10^3$. ΔCE_W^U is the world-welfare range (max−min) along the unit-of-account margin on the HR edge, in $\times 10^3$. The U/D ratio $\Delta CE_W^U / \Delta CE_W^D$ is computed from unrounded welfare values. x_W^* is the HR-edge value maximizing world CE welfare. The U/D ratio is hump-shaped in adoption, peaking at $\omega = 0.35$ (274×), but remains above 50× throughout. The world-welfare-maximizing basket weight $x_W^* = 0.35$ is invariant across all specifications.

Table 4: Two-Dimensional Welfare Optima, Decomposition, and Conflict

Panel A. Two-dimensional welfare optima		
Objective	x_d^*	x_u^*
World	0.12	0.34
HF bloc	0.00	0.09
R bloc	0.19	0.87

Panel B. Welfare-loss decomposition at own optimum							
		Second moments ($\times 10^3$)			Welfare shares (%)		
Objective	$-CE$ ($\times 10^3$)	M_c^2	M_{gap}^2	M_π^2	s_c	s_{gap}	s_π
World	3.781	155.7	578.2	4.94	61.8	38.2	<0.1
HF bloc	0.958	56.1	46.3	2.63	87.9	12.1	<0.1
R bloc	2.961	174.7	135.8	9.00	88.5	11.5	<0.1

Panel C. Two-dimensional welfare conflict ($\omega = 0.50$)			
	R bloc	HF bloc	World
Own optimum $CE_j(x_j^*)$ ($\times 10^{-3}$)	−2.961	−0.958	−3.182
Under coordinator $CE_j(x_W^*)$ ($\times 10^{-3}$)	−6.172	−1.687	−3.182
Worst imposition $CE_j(x_{-j}^*)$ ($\times 10^{-3}$)	−9.627	−8.186	−6.445
Total loss Loss_j ($\times 10^{-3}$)	6.666	7.228	—
Coordinator gain CG_j ($\times 10^{-3}$)	3.454	6.499	3.262
Residual loss after coordination ($\times 10^{-3}$)	3.211	0.730	—
Coordinator recovery rate	52%	90%	—

Notes. All computations at $\omega = 0.50$. Parameters as in Table 1. Backing aligned to denomination throughout ($\theta^a = \theta^d$); x_d and x_u denote R-currency weights with $\theta(x) = ((1-x)/2, (1-x)/2, x)$. Grid step 0.01 over $[0,1]^2$. *Panel A:* unconstrained two-dimensional welfare maximization. *Panel B:* welfare loss and discounted second moments at each bloc's own optimum; welfare-weighted shares sum to 100. *Panel C:* welfare evaluated at own optimum, world optimum, and worst unilateral imposition; these rows are the $\omega = 0.50$ snapshot of the right panel of Figure 8.

Table 5: Welfare Non-Separability Between Design Dimensions

Panel A. Welfare gain from optimizing x_u given $\bar{x}_d$				
$\bar{x}_d$	$x_u^*(R)$	ΔCE_R^U ($\times 10^{-3}$)	$x_u^*(HF)$	ΔCE_{HF}^U ($\times 10^{-3}$)
0.00	0.88	3.361	0.09	0.651
0.17	0.87	3.308	0.09	0.646
0.25	0.87	3.272	0.09	0.642
0.33	**0.87**	**3.230**	**0.09**	**0.637**
0.42	0.86	3.175	0.09	0.631
0.50	0.85	3.120	0.09	0.625
0.67	0.84	2.986	0.09	0.610
0.87	0.81	2.804	0.10	0.590
1.00	0.80	2.676	0.10	0.575
Panel B. Welfare gain from optimizing x_d given $\bar{x}_u$				
$\bar{x}_u$	$x_d^*(R)$	ΔCE_R^D ($\times 10^{-3}$)	$x_d^*(HF)$	ΔCE_{HF}^D ($\times 10^{-3}$)
0.00	0.00	0.285	0.00	0.097
0.17	1.00	0.114	0.00	0.088
0.25	1.00	0.294	0.00	0.085
0.33	**1.00**	**0.369**	**0.00**	**0.076**
0.42	1.00	0.362	0.00	0.047
0.50	1.00	0.300	0.19	0.003
0.67	0.94	0.070	1.00	0.303
0.87	0.19	0.007	1.00	1.057
1.00	0.14	0.020	1.00	1.732

Notes. All computations at $\omega = 0.50$. Parameters as in Table 1. *Panel A:* for each fixed denomination $\bar{x}_d$, the welfare gain from optimizing x_u relative to the knife-edge UoA benchmark $CE_j(\bar{x}_d, 1/3)$. *Panel B:* for each fixed unit of account $\bar{x}_u$, the welfare gain from optimizing x_d relative to the knife-edge benchmark $CE_j(1/3, \bar{x}_u)$. **Bold rows**: knife-edge benchmark in each panel.