Stablecoin flows and spillovers to FX markets[1]

Iñaki Aldasoro BIS
Paula Beltran IMF
Federico Grinberg IMF

March 26, 2026

Abstract

Using data on four USD-pegged stablecoins and 27 fiat currencies, this paper documents spillovers from stablecoin-based foreign exchange (FX) to traditional FX markets. We document a gap between the cost of acquiring dollars via stablecoins and via the spot FX market (parity deviations). To establish a causal link between stablecoin flows and FX markets, we use a granular instrumental variable that exploits idiosyncratic shocks to stablecoin net inflows in other currencies. Our estimates indicate that a 1% exogenous increase in net stablecoin inflows raises parity deviations by 40 basis points, depreciates the local currency, and widens the dollar premium in synthetic funding markets (covered interest parity (CIP) deviations). A model of constrained arbitrage rationalizes these findings and provides structural foundations for the identification strategy. Counterfactual simulations show that halving cross-market frictions would attenuate CIP spillovers by roughly one-half and cut exchange rate effects by nearly one-third. A dynamic extension that closely matches the empirical impulse responses shows that spillovers grow disproportionately when intermediaries suffer losses, as depleted capital reduces their capacity to absorb further shocks. Our results establish stablecoins as an emerging segment of global currency markets with direct implications for financial stability.

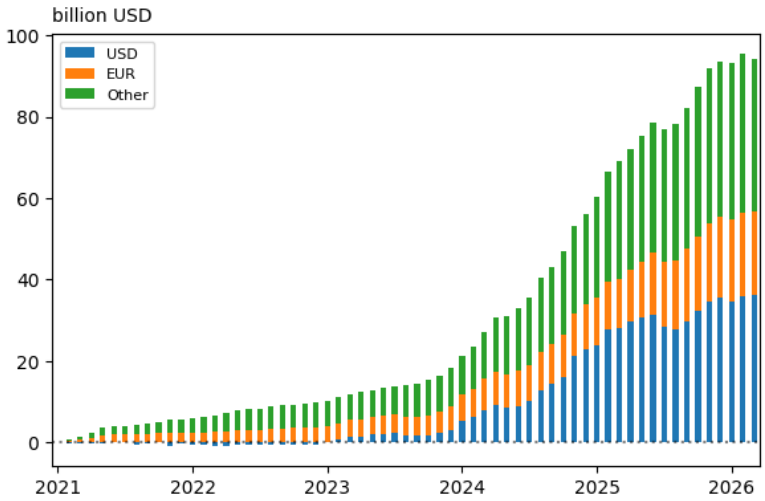
Keywords: Stablecoins, foreign exchange, market segmentation, capital flows, arbitrage.
JEL Classification: F31, G15, G12, G23, F38.

[1]Aldasoro is with the Bank for International Settlements (BIS). Corresponding author: inaki.aldasoro@bis.org. Beltran and Grinberg are with the International Monetary Fund (IMF). The views expressed herein are those of the authors and should not be attributed to the BIS, the IMF, the IMF Executive Board, or IMF management. We thank Hiroyuki Kubota for outstanding support with the data in the early stages of this project. The paper has benefited from insightful comments by Tobias Adrian, Chris Erceg, Jon Frost, Marcello Miccoli, Andrew Usher, seminar participants at the IMF and participants at the thematic group meeting on financial innovation at the BIS. All remaining errors are our own. Declarations of interest: none.

1 Introduction

Stablecoin markets have experienced rapid growth and their potential impact on the financial system has placed them at the center of policy discussions. Stablecoins are overwhelmingly denominated in United States dollars (USD). But they also have global reach: Figure 1 shows that more than 70 percent of cumulative net inflows from fiat currencies into stablecoins originate from non-USD currencies. This implies that a majority of stablecoin transactions inherently involve a foreign exchange (FX) conversion relative to USD, creating a parallel, stablecoin-based FX ecosystem. This observation raises important questions: how integrated are these stablecoin-based FX markets with their traditional, off-chain counterparts, and what are the consequences of this new venue for traditional FX markets and global capital flows?

Figure 1: Cumulative net inflows from fiat currencies to stablecoins



Note: This figure shows the cumulative sum of daily net inflows from fiat currencies into USD-pegged stablecoins (USDT, USDC, DAI, and BUSD), measured in billions of USD. The sample covers all exchanges and fiat currencies available in our dataset from 2021 to 2025.
Source: Authors' calculations based on Crypto Compare exchange-level trading data.

This paper provides a systematic analysis of the role of stablecoins in global currency markets. We make four contributions. First, using daily data on four USD-pegged stablecoins traded against 27 fiat currencies, we document substantial price gaps between acquiring dollar exposure via stablecoins and via traditional FX markets — what we term parity deviations — and show that these co-move with stablecoin flows and local currency depreciation. Second, we develop a model of segmented arbitrage in which a globally fixed stablecoin supply is allocated across countries by balance-sheet-constrained intermediaries, generating spillovers from stablecoins to covered interest parity (CIP) deviations and exchange rates, and motivating the instrument we use for causal identification. Third, using a granular instrumental variable that exploits idiosyncratic flow shocks to stablecoin net inflows in other currencies, we establish that a one percent exogenous increase in net stablecoin inflows raises parity deviations by approximately 40 basis points (bp), depreciates the local currency by 5 bp, and widens the short-term dollar

premium by 5-10 bp. Fourth, counterfactual exercises show that cross-market frictions are the primary driver of spillovers. A dynamic extension with intermediary wealth dynamics matches the empirical impulse responses while revealing scope for nonlinear amplification during stress.

Understanding these mechanisms requires data on where and how stablecoins interface with fiat currencies. Conversion between stablecoins and fiat currencies occurs primarily on centralized exchanges (CEXs) and regulated on- and off-ramp providers (Azar et al., 2022; Baughman et al., 2022).[2] These platforms serve as the principal interface between traditional money systems and stablecoin markets, and form a new type of parallel FX ecosystem that differs from traditional markets in its participant base, regulatory oversight and susceptibility to official intervention.

We build a novel dataset from 64 CEXs for the four major USD-pegged stablecoins traded against 27 fiat currencies. Our daily data span January 2021 to November 2025 and include both prices and flows. As a new parallel ecosystem that interfaces with traditional FX markets, stablecoins give rise to no-arbitrage conditions similar to traditional financial markets. Consider an agent in an emerging market seeking to acquire a stablecoin using the local currency: she can do this directly by using the local currency to purchase the stablecoin on a CEX at a price in units of local currency per stablecoin; or alternatively she can do this indirectly by first using the local currency to buy dollars at the spot FX rate and then using the dollars to buy the stablecoin at a price quoted in dollars per stablecoin. In a frictionless and fully integrated market, the law of one price dictates that the cost of the two routes should be equivalent. When these two routes yield different prices, we refer to the existence of parity deviations.

We start by providing evidence of substantial parity deviations for the currencies in our sample. For some currencies, these deviations average several percentage points and exhibit significant volatility. These parity deviations from no-arbitrage conditions are systematically larger for currencies of economies experiencing macroeconomic instability, such as high inflation, or those that employ capital flow management measures. This pattern suggests that the crypto-based FX market may provide a distinct venue for currency transactions that reflects local funding stresses and offers an alternative channel for capital movements. The existence of these sizable and heterogeneous deviations is the first of our three stylized facts, pointing to frictions in arbitrage between crypto and traditional FX venues.

These parity deviations, as well as traditional FX rates and no arbitrage conditions such as CIP, correlate with currency-specific stablecoin flows. For each stablecoin-fiat currency pair and each period t we define a measure of net inflows that captures net new funding from a

[2]Recent industry market data indicate that the vast majority of observable stablecoin trading occurs on centralized trading venues: one aggregated liquidity analysis finds that only around 5% of stablecoin trades are executed on decentralized exchanges, with the remaining $\tilde{9}5\%$ falling on centralized exchanges and associated order books. This figure is a proxy for venue concentration in stablecoin markets (and does not isolate strictly stablecoin-to-fiat pairs, which depend on proprietary data on fiat rails), but it illustrates the dominance of centralized platforms in facilitating stablecoin trades. See Kaiko Research (2025).

specific fiat currency into the stablecoin ecosystem. In our second and third stylized facts, we show that large net inflows are linked with larger parity deviations and local currency FX depreciation. These stylized facts underscore the presence of frictions to arbitrage and seem to validate policy-makers concerns, especially in emerging markets, about the potential impact of stablecoins in their domestic FX markets.

A central challenge in establishing a causal link from stablecoin flows to parity deviations and traditional FX markets is endogeneity. For example, an anticipated depreciation of a local currency could simultaneously drive capital flight into stablecoins (increasing measured inflows) and cause the depreciation and parity deviations themselves. This reverse causality, alongside potential omitted variables such as unobserved global risk factors, confounds a naïve regression of parity deviations on stablecoin flows. Consequently, the correlated relationships we document as stylized facts cannot be directly interpreted as causal.

We develop a theoretical model to rationalize these findings through a unified framework. We then provide a structural instrument to identify causal effects from stablecoin flow shocks to parity deviations. The model features home households who demand both stablecoins and synthetic dollars via FX swaps, which serve different functions and are therefore imperfect substitutes. A balance-sheet-constrained intermediary connects the two markets, facing costs that depend on total currency exposure. Stablecoin issuers' supply is treated as perfectly elastic at the dollar peg. Cross-country linkages arise instead from cross-book traders active on multiple fiat-stablecoin exchange books, who allocate a finite conversion budget across countries. This reallocation channel creates the cross-country linkages that justify our granular instrument. When conversion urgency declines in country A, cross-book participants redirect activity to country B. As such, demand for stablecoins in Country B increases and stablecoin parity deviation rise even if local demand conditions are unchanged.

Beyond the stablecoin parity deviation, additional channels operate through the intermediary's balance sheet. Expanding stablecoin positions raises the marginal cost of swap provision, lowering CIP deviations (increasing the dollar premium). Through standard channels, lower CIP deviations (more expensive synthetic dollars) reduce synthetic dollar borrowing and increase spot dollar purchases, depreciating the local currency. The model thus endogenously generates the precise triad of effects we estimate empirically: higher stablecoin parity deviations, lower CIP deviations (a higher dollar premium), and local currency depreciation.

Crucially, the model also provides the structural foundation for our identification strategy. We show that ordinary least squares (OLS) regressions of prices on stablecoin flows are biased due to correlation between local demand and unobserved factors (both global shocks that affect all markets and local confounders). The granular instrumental variable (GIV) strategy, using demand shocks in other currencies as instruments, identifies the causal effects because these shocks affect local outcomes through reduced relative supply but are uncorrelated with local demand conditions. We characterize exactly which structural parameters are identified

by instrumented regressions, including the marginal cost of intermediation, the cross-market spillover intensity, and the exchange rate sensitivity to stablecoin flows. The ratio of the estimated coefficients for deviations in CIP and stablecoin parity provides a direct test of market integration: values close to zero indicate segmented markets, while values close to one indicate highly integrated markets where stablecoin stress transmits strongly to traditional funding conditions.

Building on these insights, we construct a granular instrumental variable from idiosyncratic flow shocks to stablecoin-fiat pairs in other currencies, purged of common global factors. The IV estimates confirm the model's predictions: a one percent shock to net stablecoin inflows raises parity deviations (by approximately 40 basis points), depreciates the local currency (by 5 basis points), and widens the short-term dollar premium (by 5-10 basis points), providing evidence of causal spillovers from the crypto ecosystem to traditional foreign exchange markets.

Beyond providing a rationale for our causal identification, the structural model enables counterfactual analyses of alternative market configurations. We find that cross-market frictions (the cost that arbitrageurs face when holding mismatched positions across stablecoins and swaps) are the primary driver of spillovers to traditional FX markets. Halving these costs would reduce stablecoin parity deviation responses by around 10 percent, attenuate CIP spillovers by roughly half, and cut exchange rate depreciation effects by one-third. In contrast, lowering stablecoin-specific costs primarily compresses parity deviations, with limited pass-through to CIP or spot rates.

A dynamic extension of the model, in which intermediary risk-bearing capacity varies with wealth, generates impulse responses whose shape and persistence closely match our empirical estimates. The dynamic framework reveals an additional channel: adverse shocks deplete intermediary capital, reducing risk-bearing capacity and amplifying subsequent price responses. When flow shocks coincide with redemption frictions — as during episodes such as the May 2022 Terra/Luna collapse — this amplification can multiply baseline spillover effects several-fold, underscoring that spillovers can escalate sharply during stress episodes.

These counterfactuals reveal a policy-relevant asymmetry. Interventions that ease stablecoin market access, such as regulatory clarity or improved on-ramp infrastructure, would reduce parity deviations without destabilizing traditional FX markets, whereas tightening of arbitrageur constraints across all dollar funding venues could amplify cross-market transmission. The counterfactuals also speak to the consequences of stablecoin market growth: as demand elasticity rises with market maturation, the same flow shocks generate larger price responses, suggesting that spillover risks may increase as the market develops.

Our findings demonstrate that stablecoin markets are already linked to traditional finance, with spillovers that affect currency stability and funding conditions. This has direct implications for policymakers concerned with monetary policy autonomy and financial stability, particularly in emerging markets where these effects are most pronounced. Looking ahead, our work opens

several avenues for future research, including a deeper investigation of the heterogeneity of these effects across country characteristics, the role of specific stablecoin governance structures, and the potential for these new markets to alter the international transmission of financial shocks.

Related literature. Our paper connects to four strands of literature. First, work on CIP deviations and segmented FX markets (Sushko et al., 2016; Du et al., 2018; Rime et al., 2022), understood theoretically through limits to arbitrage (Shleifer and Vishny, 1997) and the risk-bearing capacity of specialized intermediaries (Gabaix and Maggiori, 2015; Itskhoki and Mukhin, 2021). We embed a new, crypto-based segment into this market structure and show that it generates measurable spillovers to traditional FX parity conditions.

Second, research on arbitrage in cryptocurrency markets, including cross-exchange price disparities (Makarov and Schoar, 2020), triangular arbitrage (Franz and Valentin, 2020), crypto carry trades (Schmeling et al., 2023), and crypto-facilitated capital flows (Graf von Luckner et al., 2024). We extend this work by systematically quantifying the price gaps between stablecoins and traditional FX venues and establishing their causal effects on traditional markets.

Third, a rapidly expanding literature on stablecoins covering stability (D'Avernas et al., 2023; Lyons and Viswanath-Natraj, 2023), adoption (Bertsch, 2023), runs (Ahmed et al., 2025; Gorton et al., 2026), regulation (Goel et al., 2025), market structure (Ma et al., 2023), capital flows (Reuter, 2025; Auer et al., 2025),[3] effects on commercial paper markets (Barthelemy et al., 2023; Kim, 2025a), and Treasury yields (Ahmed and Aldasoro, 2025; Kim, 2025b). Relatedly, Gorton et al. (2025) study how stablecoins develop a convenience yield, documenting that most stablecoins carry a negative convenience yield and identifying aggregate conditions, technology, reputation, and dollar demand as key drivers of stablecoins' distance to "no-questions-asked" status. We contribute by providing the first structural framework linking stablecoin markets to traditional FX outcomes.

Fourth, our identification strategy builds on the GIV approach of Gabaix and Koijen (2024), which we apply to the multi-currency stablecoin market to isolate exogenous variation from idiosyncratic,high-frequency shocks in other currency pairs.

Roadmap. The rest of the paper is organized as follows. Section 2 introduces the main data sources used in the paper and lays out the key definitions needed. Section 3 presents three stylized facts we uncover for the link between stablecoins and FX markets and briefly discusses identification challenges. In Section 4 we present a model of segmented arbitrage that provides a structural identification framework and derives our strategy for identifying causal effects. Section 5 presents the empirical specification, our empirical approach, our main results, and their structural interpretation. Section 6 presents counterfactuals grounded on the model's

[3] See also Graf von Luckner et al. (2023) for related work on bitcoin and capital flows.

estimated parameters, including a dynamic extension that characterizes impulse responses and balance-sheet amplification. Finally, Section 7 concludes.

2 Data and definitions

Data. Our analysis leverages a novel dataset to study the intersection of crypto-based and traditional FX markets. The primary data source is CryptoCompare, from which we collect daily, exchange-level data from the secondary market for cryptoassets (including stablecoins). Our sample encompasses 64 centralized exchanges (CEXs), including the largest and most liquid venues such as Binance, Coinbase, and Kraken.

We focus on four major USD-pegged stablecoins (USDT, USDC, DAI, and BUSD) that are traded against the broad set of 27 fiat currencies available in our dataset. For each stablecoin-fiat currency pair on each exchange, we collect daily closing prices, trading volumes, and gross flows (inflows from fiat to stablecoin and outflows from stablecoin to fiat). To construct currency-level aggregates, we calculate volume-weighted averages of prices and sum flows across all exchanges where a given stablecoin-fiat pair is traded. Our sample period spans from January 2021 to November 2025, capturing a period of significant growth and volatility in both crypto and traditional finance.

This granular data structure allows us to observe a parallel, stablecoin-based FX ecosystem. A key institutional feature is that stablecoin inflows from non-USD currencies embed an implicit FX conversion. The resulting parallel FX ecosystem stands in contrast to traditional, "off-chain" FX markets, which are characterized by different participants, regulatory oversight, and policy interventions (including applying capital flow management measures).

Table 1 provides descriptive statistics for USDT, the dominant stablecoin in our sample, traded against 27 fiat currencies. For each fiat currency listed in the table, we report the number of exchanges (N), the average daily trading *volume* in millions of USD, the annualized price *volatility*, and summary statistics for the parity deviations (defined below). The table reveals substantial heterogeneity. Trades with the USD as the quoted currency exhibit by far the largest volume and the lowest volatility. In contrast, pairs involving currencies like the Argentine peso (ARS) or Nigerian naira (NGN) show lower volumes but dramatically higher volatility and larger average parity deviations.

Definitions. We define a measure of price disparity between the stablecoin-based and traditional FX markets. Consider an agent seeking to acquire a stablecoin s using local fiat currency f. She has two routes. The *direct* route involves purchasing the stablecoin directly on a CEX at price $P_t^{s,f}$ (units of f per s). The *indirect* (synthetic) route involves first converting currency f into USD at the traditional spot FX rate $e_t^{\text{USD},f}$ (units of f per USD), and then using the USD to

Table 1: Descriptive statistics for USDT

				Stablecoin parity deviations							Stablecoin parity deviations		
Quote	*N*	Volume	Volatility	Mean	Median	Max	Quote	*N*	Volume	Volatility	Mean	Median	Max
USD	20	333.13	0.76%	0.05%	0.03%	0.80%	PHP	1	1.45	5.21%	0.27%	0.22%	1.80%
KRW	4	123.73	11.50%	2.51%	1.90%	10.53%	UAH	2	1.18	9.59%	1.67%	0.85%	10.28%
EUR	15	123.27	10.81%	0.31%	0.23%	2.21%	COP	2	1.15	12.25%	1.49%	1.44%	5.56%
TRY	6	77.07	6.28%	1.08%	0.68%	20.98%	CAD	2	0.89	7.27%	0.34%	0.25%	8.37%
BRL	10	16.02	13.80%	0.81%	0.71%	3.80%	PLN	3	0.72	13.92%	2.32%	0.64%	40.94%
THB	1	14.16	11.02%	0.77%	0.43%	8.50%	AED	1	0.71	1.82%	0.14%	0.12%	0.92%
GBP	8	12.24	9.16%	0.30%	0.22%	2.99%	INR	3	0.63	14.09%	4.83%	4.78%	13.19%
MXN	2	7.48	11.36%	0.48%	0.37%	2.97%	ZAR	2	0.23	12.06%	2.05%	1.96%	11.38%
TWD	1	5.60	15.60%	0.52%	0.43%	2.52%	SGD	2	0.12	6.68%	1.59%	0.26%	24.11%
IDR	2	3.56	8.55%	0.36%	0.28%	2.70%	JPY	2	0.09	29.12%	0.41%	0.23%	13.81%
ARS	2	2.59	23.10%	1.06%	0.61%	14.80%	NZD	1	0.05	15.63%	1.59%	0.46%	16.51%
CHF	1	2.27	14.68%	0.31%	0.23%	3.35%	NGN	2	0.04	13.68%	4.54%	3.44%	24.39%
HKD	1	2.20	2.26%	0.20%	0.15%	1.26%	KES	1	0.00	83.38%	3.25%	1.26%	13.74%
AUD	5	2.07	14.45%	0.40%	0.29%	8.05%							

Note: N denotes the number of exchanges included in the sample. *Volume* refers to daily trading volume measured in million USD, averaged over 2025Q2. *Volatility* is reported for 2025Q2, expressed at an annualized percentage rate. The mean, median, and maximum of parity deviations are computed using observations after 2021, following equation (1). All series are constructed from exchange-level data. Parity deviations for ARS, NGN, and UAH are measured based on parallel rates estimated by Graf von Luckner et al. (2024).
Source: Authors' calculations based on Crypto Compare exchange-level trading data.

purchase stablecoin s at price $P_t^{s,\text{USD}}$. In a frictionless, fully integrated market, the law of one price should hold, and the cost of both routes should be identical. We formalize the deviation from this parity as:

$$D_t^{s,f} \equiv \frac{P_t^{s,f}}{P_t^{s,\text{USD}} e_t^{\text{USD},f}} \times 100 \tag{1}$$

which compares the cost of the direct purchase (numerator) to the cost of the indirect purchase (denominator). The interpretation is as follows: By construction, $D_t^{s,f} = 1$ implies that there are no arbitrage opportunities between the two routes. That is, the stablecoin commands the same price regardless of the purchase route. If instead $D_t^{s,f} > 1$, then stablecoin s is relatively expensive in terms of fiat currency f relative to the synthetic USD route, suggesting an arbitrage strategy of buying stablecoin s in USD and selling them in currency f. Conversely, if $D_t^{s,f} < 1$, s the stablecoin is relatively cheap in local currency f, favoring the opposite arbitrage direction (buy in f, sell in USD).

We compute this parity deviation for all stablecoin-fiat pairs in our sample. Specifically, for each fiat currency-stablecoin pair, we calculate the deviation measure $D_t^{s,f}$ as defined above, and then summarize these deviations across countries. In our sample, the Argentine Peso (ARS), Nigerian Naira (NGN), and Ukranian Hryvnia (UAH) had periods with active on-shore parallel exchange rate markets that deviated from the official rate. For these currencies we use parallel exchange rates as reported by Haver and IMF country desks to capture effective trading

conditions, as these are more representative for agents without access to a rationed official exchange rate.[4] In Table 1, we also report descriptive statistics based on country-level average parity deviations, where each observation is weighted by trading volume in the corresponding fiat-stablecoin market.

To capture funding dynamics, for each fiat–stablecoin pair (f, s) and period t, we define the net inflow rate as:

$$g_{f,s,t} = \frac{\text{inflows}_{f,s,t} - \text{outflows}_{f,s,t}}{\text{market size}_{f,s,t-1}}, \tag{2}$$

where $\text{inflows}_{f,s,t}$ and $\text{outflows}_{f,s,t}$ are the total volumes flowing between fiat f and stablecoin s across the relevant exchanges, and $\text{market size}_{f,s,t-1}$ is outstanding stock of that stablecoin-fiat pair at the end of the previous period. This measure is expressed in percentage terms and winsorized at the 1/99th percentile to mitigate the influence of outliers. It represents net new funding from a specific fiat currency into the stablecoin ecosystem, i.e. net inflows are recorded when fiat currency is exchanged for stablecoins – equivalent to fiat-to-stablecoin conversion.

To connect our analysis to the extensive literature on frictions in traditional FX markets, we construct measures of CIP deviations. CIP is a fundamental no-arbitrage condition linking interest rates and spot and forward exchange rates. The CIP deviation, δ, measures the cost of obtaining dollars synthetically (via FX swaps) relative to borrowing dollars directly. When $\delta < 0$, synthetic dollar funding is more expensive than direct dollar borrowing; in other words, there is a dollar premium in the swap market.[5]

As documented by Du et al. (2018), CIP deviations were negligible before the Great Financial Crisis but have become a persistent feature of FX markets since then, reflecting factors such as bank balance sheet costs and market segmentation. In the context of our paper, we investigate whether shocks originating in the stablecoin market spill over and affect these traditional funding arbitrage conditions. We construct this measure for 25 currencies in our sample at 3-month and 12-month maturities.

3 Stylized facts

Armed with our key variable definitions, we now document three novel stylized facts that characterize the relationship between stablecoin flows and traditional FX markets. These facts provide the initial, reduced-form evidence of a link between the crypto-based and traditional

[4] This is also a conservative stance that reduces the magnitude of parity deviations.

[5] This matches the empirical definition in Section 4: $\delta \equiv (F - S)/S - (i^H - i^{USD})$, where i^H and i^{USD} are interest rates and F is the forward rate.

FX ecosystems, setting the stage for our causal identification strategy.

Stylized fact #1: *There are sizable parity deviations between stablecoin and traditional FX markets, with significant heterogeneity across currencies.*

Using the definition in equation (1), our first stylized fact is the existence of non-trivial and persistent price gaps between the direct purchase of a stablecoin in a local currency and the indirect route mediated by the USD.

Figure 2 plots the distribution of time-averaged parity deviations for USDT across the fiat currencies in our sample.The left panel shows that for a set of major and highly liquid currencies (e.g., USD, EUR, GBP), parity deviations are small and tightly distributed around zero. This suggests that arbitrage between the crypto and traditional markets for these currencies is relatively effective. For other currencies the distribution is more skewed (e.g. the Colombian peso (COP), the South African rand (ZAR), the Brazilian real (BRL), and the Turkish lira (TRY). For the currencies on the right panel, average parity deviations are considerably larger (even up to one order of magnitude) and dispersion can be quite large – hence why we plot them separately.[6] This stylized fact echoes findings in the literature on traditional FX markets (Kalemli-Özcan and Varela, 2024), extending them to the emerging stablecoin system and, as discussed later, highlighting similar underlying frictions. A related measure appears in Gorton et al. (2025), who construct a "stablecoin basis" comparing stablecoin-implied and traditional FX rates. Their focus is on how this basis reflects foreign dollar demand that makes individual stablecoins more money-like; our interest is instead in parity deviations as a price of segmentation between crypto and traditional FX venues and in their causal spillovers to CIP deviations and exchange rates.

To establish our second stylized fact, we estimate linear projection regressions that trace out the dynamic impact of net stablecoin inflows ($g_{f,s,t}$) on our variables of interest. The baseline pools all four USD-pegged stablecoins (USDT, USDC, DAI, and BUSD) traded against the fiat currencies in our sample. We estimate the following specification for horizons $h = 0, 1, ..., 10$:

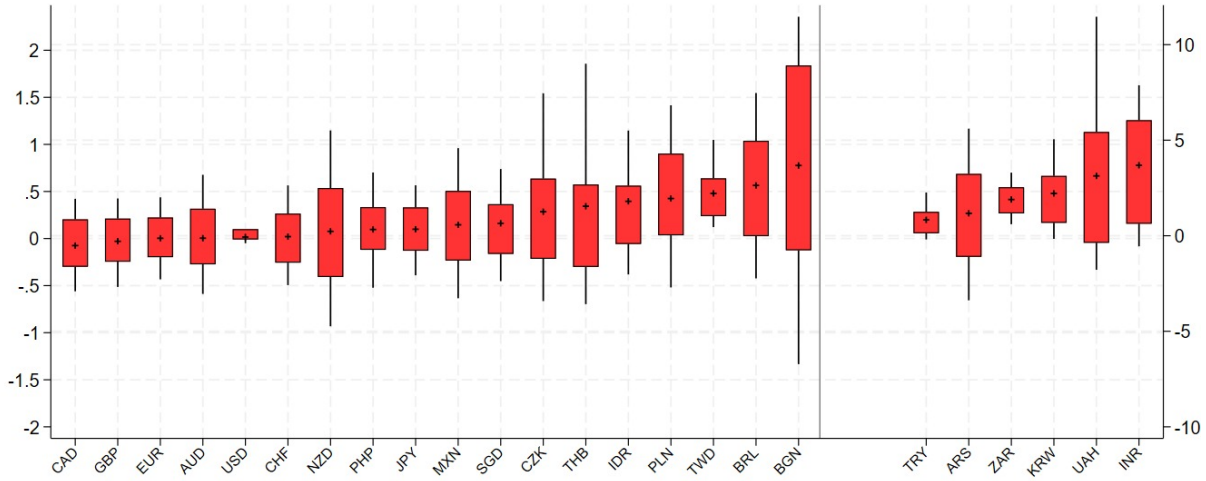
$$Y_{f,s,t+h} = \beta^h g_{f,s,t} + \gamma \mathbf{X}_{f,s,t} + \delta_{f,s} + \zeta_{f,s,t+h}, \tag{3}$$

where $Y_{f,s,t+h}$ represent our three outcome variables of interest (parity deviations, the change in exchange rates, and CIP deviations) for domestic currency f and stablecoin s at horizon h, $g_{f,s,t}$ denote net inflows into stablecoins as defined in equation (2), $\delta_{f,s}$ are stablecoin-fiat fixed effects, X captures controls and $\zeta_{f,s,t+h}$ is an error term. The coefficient β^h traces out the conditional projection of Y at horizon h to a unit inflow change.[7]

[6] The heterogeneity in the sample maybe reflecting different exchange rate regimes, degrees of capital account openness, and/or levels of macroeconomic volatility. We leave a deeper analysis of this heterogeneity for further research.

[7] We scale $g_{f,s,t}$ in percentage terms and winsorize at the 1st and 99th percentiles to mitigate the influence of outliers.

Figure 2: Distribution of parity deviations by currency (in percent)



Note: This figure shows the distribution of parity deviations against the USD by fiat currency (black median line with interquartile range boxes, with whiskers denoting the 5th–95th percentiles). The sample includes all pairs of fiat currencies and USD-pegged stablecoins between 2021 and 2025. All series are constructed from data at the fiat-stablecoin level with volume-weighted averages across exchanges. Parity deviations for the Argentinean peso (ARS), the Nigerian naira (NGN), and Ukrainian hryvnia (UAH) are measured based on parallel rates.

Our second stylized fact links net inflows into stablecoins with developments in parity deviations, exchange rates and CIP deviations:

Stylized fact #2: *Net inflows into stablecoins are correlated with: (**2a**) parity deviations, (**2b**) (non-crypto) exchange rate depreciation and (**2c**) a decline in CIP deviations (a higher dollar premium).*

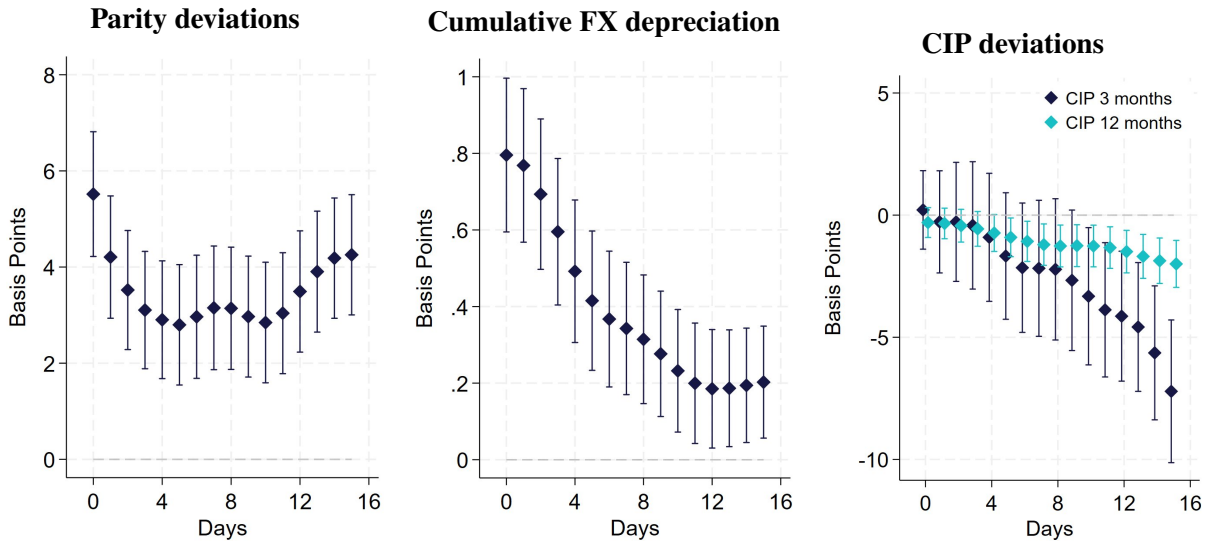
Figure 3 presents the results from estimating the linear projection regressions to support the stylized fact. Stablecoin inflows are associated with a contemporaneous increase in parity deviations (first panel). A positive increase to net stablecoin inflows is associated with an immediate and statistically significant jump in parity deviations. This is consistent with a model of segmented markets where increased demand for stablecoins in a local currency pushes up their price on crypto exchanges relative to the indirect price ($P_t^{s,USD} e_t^{USD,f}$), and where arbitrage forces are insufficient to instantly close the gap.

Perhaps more strikingly, net stablecoin inflows are associated with a subsequent depreciation of the domestic currency in the traditional spot FX market (middle panel). This provides suggestive evidence that flows into stablecoins are not confined to the crypto ecosystem but can generate tangible pressures in traditional FX markets, potentially by increasing the net supply of the local currency in the spot market as users sell local currency to buy stablecoins.

Finally, stablecoin inflows are correlated with a decline in CIP deviations, indicating a higher dollar premium. This pattern suggests that net stablecoin inflows may also generate pressures in the FX swap market, impairing the mechanism of synthetic USD creation to some extent. The effect on short-term CIP deviations is consistent with arbitrageurs facing balance sheet

constraints that are more binding at shorter horizons, a channel our model in Section 4 will formalize.

Figure 3: The effect of stablecoin inflows on parity deviations, exchange rates and CIP deviations



Note: This figure shows the cumulative impact of a one-standard deviation increase in stablecoin net inflows as defined in equation (2) (dots) and its 95% confidence interval (lines) on fiat-stablecoin parity deviation, exchange rates and CIP deviations (all in basis points) over horizons h=0,. . . ,15 days. Sample includes all four USD-pegged stablecoins.

These reduced-form correlations establish an empirical link between stablecoin activity and a range of traditional financial variables. However, interpreting these patterns causally is challenging due to endogeneity concerns. For instance, an anticipated depreciation of the local currency could simultaneously drive investors to seek refuge in stablecoins (increasing $g_{f,s,t}$) and cause the depreciation and CIP deviations. The following section presents a model that provides a structural identification framework to estimate causal effects of stablecoin inflows.

4 Structural model and identification

This section develops a model that provides structural foundations for our identification strategy and for the counterfactual exercises in Section 6. The argument proceeds in four steps.

First, in a single Home-country setting, we show that an intermediary who connects the stablecoin and FX swap markets generates exactly the triad of effects documented in Section 3: higher stablecoin demand raises parity deviations, widens CIP deviations (a higher dollar premium), and depreciates the local currency (Sections 4.1–4.2).

Second, we show that the cross-market cost parameter $\Gamma_\times$ (the cost the intermediary faces on its total currency exposure) is the sole channel through which stablecoin stress spills over to CIP deviations. When $\Gamma_\times = 0$, the two markets decouple entirely.

Third, we extend the model to multiple countries. Cross-book participants (i.e. traders active on several fiat-stablecoin exchange books) allocate a finite conversion budget across countries. Their reallocation in response to local shocks elsewhere creates the cross-country linkages needed for identification (Section 4.2).

Fourth, we derive what the IV coefficients identify in terms of structural parameters: the marginal cost of intermediation, the cross-market spillover intensity, and a direct test of market integration (Section 4.2).[8]

4.1 The model

4.1.1 Environment

Here we illustrate the main mechanisms, before extending to a more general multicountry setting that allows us to construct a structural GIV in section 4.2.

We consider an economy with only two countries: Home, with currency H, and the United States (U.S.), with currency USD. Let S denote the spot exchange rate, quoted as units of H per USD (local currency per dollar). We normalize $S = 1$ as reference. The analysis is conducted in log deviations $s = \log S$, so that $s > 0$ corresponds to depreciation and $s < 0$ corresponds to appreciation of the local currency.

Two prices are central to our analysis. The *stablecoin parity deviation* is defined as

$$y \equiv p^{sc} - s, \tag{4}$$

where $p^{sc} = \log P^{sc}$ is the log stablecoin price in currency H. When $y > 0$, acquiring dollar exposure via stablecoins costs more than via spot FX; that is, stablecoins trade at a premium.

The *CIP deviation* is defined as

$$\delta \equiv (F - S)/S - (i^H - i^{USD}), \tag{5}$$

where i^H and i^{USD} are interest rates and F is the forward rate. When $\delta < 0$, the forward premium is insufficient to compensate for the interest differential. δ measures the cost of obtaining dollars synthetically (via FX swaps) relative to borrowing dollars directly. We adopt the sign convention that $\delta < 0$ indicates a dollar premium in the swap market, i.e., synthetic dollar funding is more expensive than direct dollar borrowing.

[8]The model we present here is static. Section 6.3 extends it to incorporate persistent demand shocks and intermediary balance sheet dynamics, following He and Krishnamurthy (2013) and Gabaix and Maggiori (2015), so that adverse shocks that deplete capital generate amplified and persistent price responses.

4.1.2 Home households

Home households demand stablecoins (for on-chain dollar access) and synthetic dollars via FX swaps (for hedging and trade finance). Because these serve different functions, they are imperfect substitutes. Demand for each asset is linear in its cost:[9]

$$D^{sc} = \bar{D}^{sc} - \alpha_{sc}\, y \tag{6}$$

$$D^{syn} = \bar{D}^{syn} + \alpha_{\delta}\, \delta \tag{7}$$

where $\bar{D}^{sc}$ and $\bar{D}^{syn}$ are baseline demands, and $\alpha_{sc}, \alpha_{\delta} > 0$ are demand elasticities. The baseline demand $\bar{D}^{sc}$ captures stablecoin demand shocks (e.g. capital flight, growth in crypto commerce, or regulatory changes affecting traditional dollar access), which are the key drivers of our empirical analysis.

4.1.3 Global stablecoin supply and relative supply

Stablecoins are issued by entities that maintain reserves of dollar-denominated assets. At the daily frequency relevant for our analysis, stablecoin supply is highly elastic. We therefore set $P_t^{s,\text{USD}} = 1$.[10]

4.1.4 US dollar supply to FX swap markets

US investors supply dollars to Home FX swap markets to earn the CIP basis.[11] We model this supply as:

We model US dollar supply to the swap market as:

$$L^{US} = \bar{L} - \alpha_L\, \delta \tag{8}$$

where $\bar{L} > 0$ is baseline supply and $\alpha_L > 0$ is the supply elasticity (supply increases when the dollar premium widens).

[9] These demand functions can be derived from a household problem that maximizes convenience yields net of acquisition costs. Specifically, households maximize $V(D^{sc}, D^{syn}) - y \cdot D^{sc} + \delta \cdot D^{syn}$, where $V(\cdot)$ is an additively separable, quadratic convenience yield. Note the positive sign on δ: when $\delta < 0$ (dollar premium), households face a cost $|\delta|$ for synthetic dollars. The first-order conditions yield the demand functions in (6)–(7). See Appendix D for an alternative specification with segmented household types.

[10] We treat stablecoin supply as perfectly elastic, so $P_t^{s,\text{USD}} = 1$. Empirically, USDT/USD deviates from par by only 0.05% on average (Table 1). If supply were finitely elastic ($S^{sc} = S_0 + \kappa\mu$, $\mu = P_t^{s,\text{USD}} - 1$), a demand increase would push $P_t^{s,\text{USD}}$ above par, partially dampening the parity deviation response. This effect is isomorphic to a higher effective demand elasticity α_{sc}: the qualitative structure of the model is unchanged, and α_{sc} should be interpreted as capturing both household price sensitivity and any supply-side adjustment.

[11] This supply-side behavior is well-documented in the CIP literature. See Du et al. (2018) and Rime et al. (2022) for evidence on the role of US financial institutions in FX swap market intermediation.

4.1.5 Intermediary

A representative intermediary (capturing the aggregate activity of FX dealers, crypto market makers, and arbitrageurs) facilitates both stablecoin allocation and FX swap provision.

Stablecoin market activity. The intermediary acquires stablecoins from the global pool (where they trade near par in USD) and sells them to Home households, who pay in local currency H. The intermediary earns the stablecoin parity deviation y per unit. We denote the intermediary's stablecoin position by B^{sc}.[12]

FX swap market activity. The intermediary facilitates synthetic dollar funding for Home households by intermediating FX swaps. Concretely, when a Home firm wants synthetic dollars, the firm borrows in local currency H, the intermediary arranges the swap so that the firm delivers H spot and receives USD with a forward contract to reverse at maturity, and the intermediary stands between Home borrowers (who need dollars) and US dollar lenders (who supply dollars to earn the basis). The intermediary earns $|\delta|$ (the absolute value of the CIP deviation) as compensation for bearing balance sheet costs and counterparty risk. We denote the intermediary's swap market position by B^{syn}.

Cost structure. The intermediary faces costs that depend on positions in both markets, based on a quadratic form with convex marginal costs, where Γ_δ, Γ_{sc}, and $\Gamma_\times$ are non-negative cost parameters governing the curvature of costs in each market:

$$C(B^{syn}, B^{sc}) = \frac{\Gamma_\delta}{2}(B^{syn})^2 + \frac{\Gamma_{sc}}{2}(B^{sc})^2 + \frac{\Gamma_\times}{2}(B^{syn} + B^{sc})^2 \tag{9}$$

The first term captures swap-specific costs from balance sheet constraints (regulatory capital requirements such as the leverage ratio or internal value-at-risk type constraints), counterparty credit risk, and margin funding costs. The second term captures stablecoin-specific costs from on-ramp/off-ramp frictions, operational expenses, and issuer counterparty risk. The third term is the *cross-market cost*, depending on total currency H exposure. This term captures the key friction: when the intermediary expands its stablecoin position (acquiring local currency from households), it increases its overall exposure to currency H, raising the marginal cost of also taking on swap positions (which also involve local currency exposure). When $\Gamma_\times > 0$, expanding the stablecoin position raises the marginal cost of swap provision, generating spillovers between markets.

[12] The stablecoin inflows defined in equation (2) map into B^{sc}. When $B^{sc} > 0$, the intermediary sells stablecoins to households, receiving local local currency in exchange.

The intermediary maximizes profits:

$$\max_{B^{syn},B^{sc}} (-\delta) \cdot B^{syn} + y \cdot B^{sc} - C(B^{syn}, B^{sc}) \tag{10}$$

Note that since $\delta < 0$ in equilibrium (dollar premium), the intermediary's revenue from swap intermediation is $(-\delta) \cdot B^{syn} = |\delta| \cdot B^{syn} > 0$.

4.1.6 Market clearing

Stablecoin market. The intermediary absorbs net demand on the Home exchange book:

$$B^{sc} = D^{sc} + S^{sc} \tag{11}$$

where $D^{sc} = \bar{D}^{sc} - \alpha_{sc}\, y$ is household demand and S^{sc} represents additional demand from other market participants.[13] Defining net demand pressure as $\bar{\rho}^{sc} \equiv \bar{D}^{sc} + S^{sc}$, we obtain $B^{sc} = \bar{\rho}^{sc} - \alpha_{sc}$. A household demand shock (higher $\bar{D}^{sc}$) or an increase in other participants' demand (higher S^{sc}) both raise $\bar{\rho}^{sc}$, increasing the quantity the intermediary must absorb and the premium it requires.

FX swap market. Home demand for synthetic dollars equals US supply plus intermediary provision:

$$D^{syn} = L^{US} + B^{syn} \tag{12}$$

Defining $\bar{\rho}^{syn} \equiv \bar{D}^{syn} - \bar{L}$ and $\tilde{\alpha}_\delta \equiv \alpha_\delta + \alpha_L$, we obtain $B^{syn} = \bar{\rho}^{syn} + \tilde{\alpha}_\delta\, \delta$

Spot market. The spot exchange rate, Q^{spot}, is determined by:

$$Q^{spot} = \bar{Q} - \eta_\delta\, \delta + \eta_y\, y - \nu_s\, s = 0 \tag{13}$$

where $\bar{Q}$ captures baseline net dollar demand in the spot market arising from fundamental factors that are independent of stablecoin and swap market conditions, $\eta_\delta, \eta_y > 0$ capture substitution toward spot when swaps or stablecoins become expensive, and $\nu_s > 0$ reflects standard exchange rate elasticities. Note the sign on η_δ: when δ becomes more negative (dollar premium widens), households substitute toward spot dollar purchases, increasing Q^{spot}.[14]

[13] This term captures demand from participants whose behavior we do not micro-found in the single-country model. For instance, traders active on multiple exchange books, over-the-counter (OTC) desks, or agents using stablecoins for remittances. In Section 4.2 we provide structure for S^{sc}, where cross-book participants allocate a finite conversion budget across markets.

[14] A more structural specification links spot market clearing directly to intermediary positions: $Q^{spot} = \bar{Q} + \beta_{syn} B^{syn} - \beta_{sc} B^{sc} + \nu_s s = 0$, where β_{syn} and β_{sc} capture the spot market flows generated by swap and stablecoin intermediation respectively. When the intermediary expands B^{sc}, it receives currency H from households and may hedge this exposure by selling H in the spot market; similarly, swap positions have spot legs that affect market clearing. Using the market clearing conditions $B^{sc} = \bar{\rho}^{sc} - \alpha_{sc} y$ and $B^{syn} = \bar{\rho}^{syn} + \tilde{\alpha}_{syn}\delta$,

4.1.7 Equilibrium and comparative statics

Intermediary optimality. The first-order conditions for the intermediary equate marginal revenue to marginal cost:

$$-\delta = (\Gamma_\delta + \Gamma_\times) B^{syn} + \Gamma_\times B^{sc} \tag{14}$$

$$y = \Gamma_\times B^{syn} + (\Gamma_{sc} + \Gamma_\times) B^{sc} \tag{15}$$

The key feature is that each market's marginal cost depends on the position in both markets through $\Gamma_\times$. Note that the left-hand side of (14) is $-\delta > 0$ since $\delta < 0$ in equilibrium.

Equilibrium characterization. Substituting market clearing into the first-order conditions yields a linear system in $(-\delta, y)$.

Proposition 1 (Equilibrium)**.** *The unique equilibrium is characterized by:*

$$y^* = \frac{1}{\Delta}\left[\underbrace{N_y \cdot \bar{\rho}^{sc}}_{\text{stablecoin demand pressure}} + \underbrace{\Gamma_\times \cdot \bar{\rho}^{syn}}_{\text{swap spillover}}\right] \tag{16}$$

$$-\delta^* = \frac{1}{\Delta}\left[\underbrace{N_\delta \cdot \bar{\rho}^{syn}}_{\text{swap demand pressure}} + \underbrace{\Gamma_\times \cdot \bar{\rho}^{sc}}_{\text{stablecoin spillover}}\right] \tag{17}$$

$$s^* = \frac{1}{\nu_s}\left(-\eta_\delta \delta^* + \eta_y y^* + \bar{Q}\right) \tag{18}$$

where $A_\delta \equiv 1 + \tilde{\alpha}_\delta(\Gamma_\delta + \Gamma_\times)$, $A_{sc} \equiv 1 + \alpha_{sc}(\Gamma_{sc} + \Gamma_\times)$, $B \equiv \alpha_{sc}\Gamma_\times$, $C \equiv \tilde{\alpha}_\delta \Gamma_\times$, $N_y \equiv A_\delta(\Gamma_{sc} + \Gamma_\times) - \tilde{\alpha}_\delta \Gamma_\times^2$, $N_\delta \equiv A_{sc}(\Gamma_\delta + \Gamma_\times) - \alpha_{sc}\Gamma_\times^2$, *and* $\Delta \equiv A_\delta A_{sc} - BC > 0$.

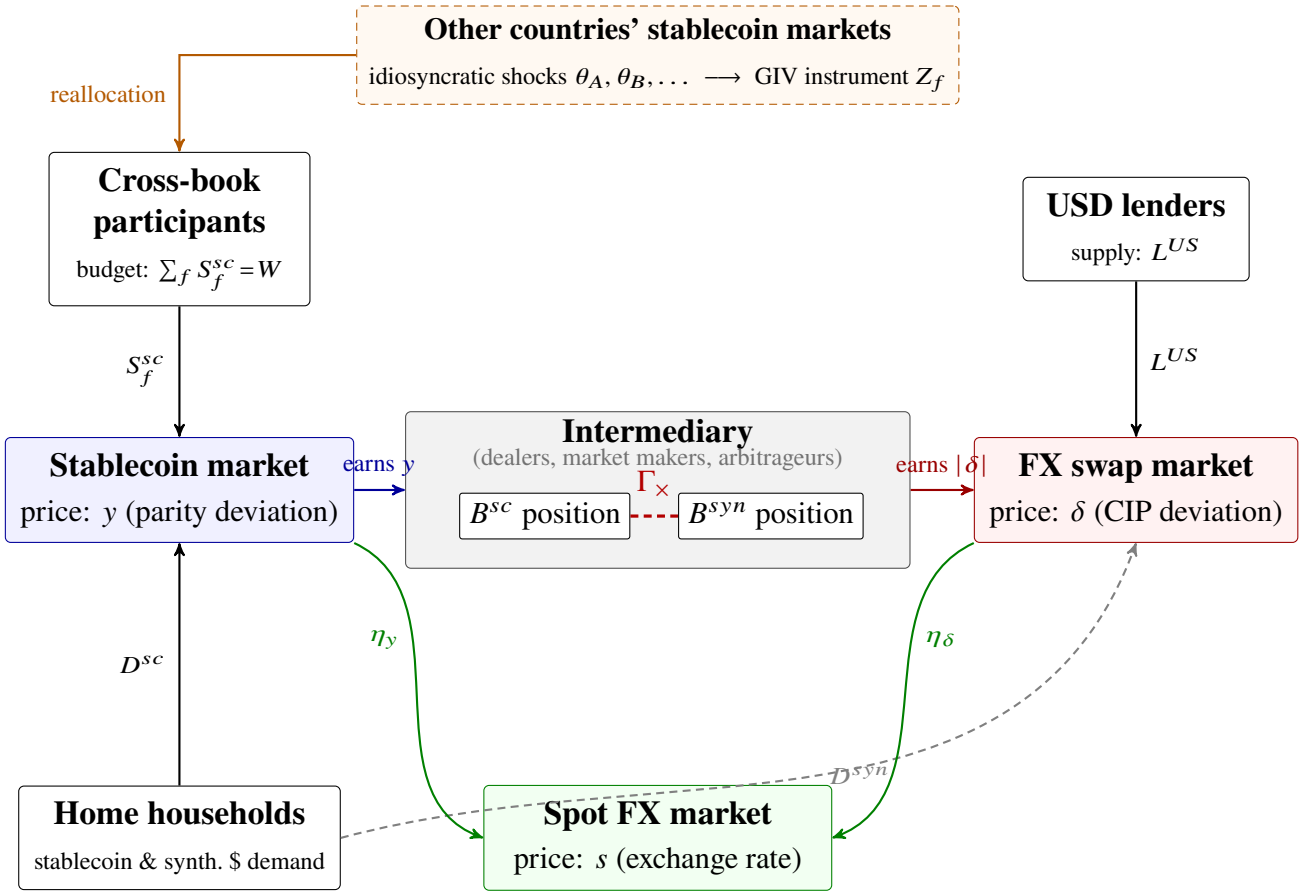
Proof. See Appendix C.

In the equilibrium, each price is driven by own-market demand pressure (weighted by N_y or N_δ, which capture own-market intermediation costs net of cross-market feedbacks) and by a spillover from the other market, weighted by the cross market spillover, $\Gamma_\times$. In particular, higher stablecoin demand raises the dollar premium through the intermediaries' balance sheet constraint,$\Gamma_\times$. When $\Gamma_\times = 0$, the two markets decouple.

Figure 4 provides a visual overview of the model's structure.

the structural specification reduces to the form in the text, with $\eta_y = \beta_{sc}\alpha_{sc}$ and $\eta_\delta = \beta_{syn}\tilde{\alpha}_{syn}$. The reduced-form coefficients thus combine the intermediary's spot market footprint with demand elasticities. We use the reduced-form approach to minimize notation while preserving the key comparative statics.

Figure 4: Schematic of the model



Note: This figure illustrates the model's structure. The intermediary connects the stablecoin and FX swap markets through two positions (B^{sc} and B^{syn}) linked by the cross-market cost $\Gamma_\times$ (red dashed line). When $\Gamma_\times > 0$, expanding the stablecoin position raises the marginal cost of swap provision, generating spillovers from stablecoins to CIP deviations and, through substitution toward the spot market (η_y, η_δ), to the exchange rate. Home households demand both stablecoins and synthetic dollars. Cross-book participants allocate a finite conversion budget ($\sum_f S_f^{sc} = W$) across countries; idiosyncratic shocks in other countries (dashed orange box) provide the identifying variation for the GIV instrument. US dollar lenders supply funds to the swap market.

Comparative statics. The equilibrium characterization reveals how stablecoin demand shocks propagate through the system. Three questions are central: does higher stablecoin demand raise parity deviations? Does it spill over to CIP? How do these effects aggregate into exchange rate movements? Proposition 2 and corollary 1 summarize the comparative statics results from our model addressing these questions.

Proposition 2 (Comparative statics)**.** *An increase in stablecoin net demand $\bar{\rho}^{sc}$:*

(i) ***raises the stablecoin parity deviation*** *(own-market effect):*

$$\frac{\partial y^*}{\partial \bar{\rho}^{sc}} = \frac{N_y}{\Delta} > 0 \tag{19}$$

(ii) ***widens the CIP deviation*** *(cross-market spillover):*

$$\frac{\partial \delta^*}{\partial \bar{\rho}^{sc}} = -\frac{\Gamma_\times}{\Delta} \leq 0 \tag{20}$$

with strict inequality when $\Gamma_\times > 0$*, and equality when* $\Gamma_\times = 0$*;*

(iii) ***depreciates the local currency*** *(total FX impact):*

$$\frac{\partial s^*}{\partial \bar{\rho}^{sc}} = \frac{1}{\nu_s}\left(-\eta_\delta \frac{\partial \delta^*}{\partial \bar{\rho}^{sc}} + \eta_y \frac{\partial y^*}{\partial \bar{\rho}^{sc}}\right) > 0 \tag{21}$$

Corollary 1 (Segmentation)**.** *When* $\Gamma_\times = 0$*, the CIP spillover vanishes and the two markets decouple* ($\partial \delta^* / \partial \bar{\rho}^{sc} = 0$)*.*

Proof. See Appendix C.

The three effects share a common origin. Higher demand forces the intermediary to expand its stablecoin position B^{sc}, raising marginal costs. The parity deviation y must rise to compensate (part i). If $\Gamma_\times > 0$, the expanded stablecoin position also raises the cost of swap provision — even with unchanged swap demand — widening the CIP deviation through a balance sheet externality (part ii). Both channels feed into spot depreciation: households substitute toward spot dollar purchases when stablecoins become expensive (weighted by η_y) and when synthetic funding costs rise (weighted by η_δ), with both effects working in the same direction (part iii). Importantly, stablecoin market stress can thereby affect exchange rates even for agents who do not participate in crypto markets, since the CIP channel raises funding costs for any firm relying on FX swaps for trade finance or hedging.

Finding a non-zero CIP spillover in the data constitutes direct evidence of balance sheet linkages ($\Gamma_\times > 0$) between stablecoin and traditional FX markets. Our empirical strategy is designed precisely to test this prediction.

4.2 Multiple currencies and identification

We now extend the model to F home countries to study cross-country spillovers and provide structural foundations for our identification strategy. The extension serves two purposes. First, it formalizes how stablecoin demand shocks transmit across countries, which underpins our identification strategy. Second, it allows us to be explicit about the sources of endogeneity that bias OLS estimates and to show how our granular instrumental variable overcomes these biases.

Multi-currency setup. Consider F home countries indexed by $f \in \{1, \ldots, F\}$. Each country has local households with stablecoin demand $D_f^{sc} = \bar{D}_f^{sc} - \alpha_{sc}\, y_f$ and an intermediary facing the cost structure from Section 4.1.5. Within each country, equilibrium is characterized by

Proposition 1, with prices determined by net demand pressure $\bar{\rho}_f^{sc} = \bar{D}_f^{sc} + S_f^{sc}$.

We now give structure to the additional demand S_f^{sc} introduced in Section 4.1.6. These participants are *cross-book participants*: traders and entities active across multiple fiat-stablecoin exchange books. A crypto-native trader on Binance with verified accounts for BRL/USDT, TRY/USDT, and KRW/USDT is a canonical example. These agents hold fiat balances in several currencies and choose how much to convert on each book in each period.

Crucially, cross-book participants face a finite total conversion budget:

$$\sum_{f=1}^{F} S_f^{sc} = W \tag{22}$$

reflecting operational constraints such as exchange limits, know-your-customer (KYC) processing bandwidth, fiat settlement capacity, and total fiat holdings. This budget constraint is the source of cross-country linkages in stablecoin markets.

Cross-book participant behavior. Cross-book participants allocate their budget based on local conditions across markets. We summarize the attractiveness of converting currency f by a *conversion urgency* parameter $\theta_f > 0$, which captures expected depreciation, regulatory risk, banking disruptions, and capital control tightening. Appendix H derives the optimal allocation from a portfolio-choice problem. The key result is:

$$S_f^{sc} = \bar{S}_f^{sc} + v_f, \qquad \sum_f v_f = 0 \tag{23}$$

where v_f is the idiosyncratic component of reallocation, driven by urgency shocks elsewhere. Cross-book participants convert more from currencies with higher urgency, and the zero-sum property follows from the budget constraint: higher conversion in one country mechanically reduces conversion in others.

Cross-country transmission. When cross-book participants redirect conversion toward country A's book (S_A^{sc} rises due to higher θ_A), their participation on other books must decline (S_B^{sc} falls for $B \neq A$). From country B's perspective, demand pressure falls ($\bar{\rho}_B^{sc}$ declines) even though local household demand $\bar{D}_B^{sc}$ is unchanged.

Proposition 3 (Cross-country spillovers)**.** *A reallocation of cross-book participants from country B to country A (S_A^{sc} increases, S_B^{sc} decreases, $\sum_f S_f^{sc} = W$) affects country B:*

$$\frac{\partial y_B}{\partial S_B^{sc}} > 0, \quad \frac{\partial \delta_B}{\partial S_B^{sc}} < 0, \quad \frac{\partial s_B}{\partial S_B^{sc}} > 0 \quad \text{for } A \neq B \tag{24}$$

A reduction in cross-book participant demand on country B's book lowers country B's stablecoin

parity deviation, narrows country B's CIP deviation (a lower dollar premium), and appreciates country B's currency.

Proof. See Appendix C. □

Proposition 3 establishes that demand shocks in one country spill over to affect prices in other countries. This cross-country transmission is the foundation of our identification strategy: we use demand shocks in other countries as instruments for local stablecoin flows. The proposition guarantees that these external shocks are relevant, as they move local prices through the relative supply and balance sheet channels.

4.3 An identification strategy with Granular Instrumental Variables

This subsection formalizes our identification strategy. We first show why OLS is biased, then construct the granular instrumental variable, and finally characterize what the IV coefficients identify in terms of the model's structural parameters (Proposition 4).

The model clarifies why OLS regressions of prices on stablecoin flows are biased and how the multi-currency structure can overcome this bias. Local stablecoin demand shocks are correlated with unobserved factors that independently affect parity deviations, CIP, and exchange rates — for instance, anticipated depreciation simultaneously drives capital flight into stablecoins and causes the depreciation itself. This correlation between u_f and the local error terms ε^y_f, ε^δ_f, and ε^s_f (formalized in Appendix A) renders OLS inconsistent. However, Proposition 3 establishes that demand shocks in other countries affect local prices while remaining uncorrelated with local confounders. This cross-country transmission provides the basis for a granular instrumental variable. The following proposition formalizes what such an instrument identifies in terms of the model's structural parameters. In particular, Proposition 4 shows that the IV coefficients identify key structural objects: the marginal cost of intermediation, the cross-market spillover intensity, and a direct test of market integration.

Define total demand on book f as $\tilde{D}^{sc}_f \equiv \bar{D}^{sc}_f + S^{sc}_f$, which combines household demand and cross-book participant demand. Our GIV isolates exogenous variation in $\tilde{D}^{sc}_f$ that is uncorrelated with local confounders. We decompose total demand into common and idiosyncratic components:

$$\tilde{D}^{sc}_f = \bar{D}_0 + \bar{S}^{sc}_f + \phi\mu + \underbrace{u^D_f + v_f}_{\equiv \tilde{u}_f}, \tag{25}$$

where $\tilde{u}_f$ is the idiosyncratic component of total demand on book f, combining local household shocks (u^D_f) and cross-book reallocation (v_f). We estimate $\tilde{u}_f$ as the residual from a factor model applied to observed net inflow rates (equation 31). The instrument for country

f is:

$$Z_f \equiv \sum_{f' \neq f} w_{f'} \hat{\hat{u}}_{f'}, \tag{26}$$

where $\hat{\hat{u}}_{f'}$ are factor-model residuals from total flows and $w_{f'} \equiv \bar{D}^{sc}_{f'} / \sum_{k \neq f} \bar{D}^{sc}_{k}$ are demand-weighted shares. The instrument aggregates idiosyncratic total demand shocks from other countries, excluding country f's own shock.[15]

Under reasonable assumptions on the shocks, the usual requirements for instruments are met. Assuming that the shocks to stablecoin demand that are idiosyncratic to country f are correlated with local unobservable factors affecting parity deviations, the instrument satisfies relevance ($\text{Cov}(g_f, Z_f) \neq 0$ because demand elsewhere reduces demand to f), exogeneity ($\text{Cov}(Z_f, \varepsilon_f) = 0$ because idiosyncratic shocks in other countries are uncorrelated with local unobservables), and exclusion (as Z_f affects outcomes only through the stablecoin channel). Appendix A formalizes this.

Proposition 4 (IV identification)**.** *Under Assumption 1, the IV coefficients identify the causal effects of stablecoin flows on prices:*

(i) Stablecoin parity deviation coefficient β_y^{IV} *— marginal cost of intermediation (Table 2, row (i)):*

$$\beta_y^{IV} = \frac{Cov(y_f, Z_f)}{Cov(g_f, Z_f)} = \frac{\partial y^{sc} / \partial \bar{\rho}^{sc}}{\gamma_g} = \frac{N_y}{A_\delta} \tag{27}$$

(ii) CIP spillover coefficient β_δ^{IV} *— cross-market externality (Table 2, row (ii)):*

$$\beta_\delta^{IV} = \frac{Cov(\delta_f, Z_f)}{Cov(g_f, Z_f)} = \frac{\partial \delta^{sc} / \partial \bar{\rho}^{sc}}{\gamma_g} = -\frac{\Gamma_\times}{A_\delta} \tag{28}$$

(iii) Exchange rate coefficient β_s^{IV} *— total FX impact (Table 2, row (iii)):*

$$\beta_s^{IV} = \frac{\partial s^{sc} / \partial \bar{\rho}^{sc}}{\gamma_g} = \frac{1}{\nu_s} \left[-\eta_\delta \beta_\delta^{IV} + \eta_y \beta_y^{IV} \right] = \frac{\eta_\delta \Gamma_\times + \eta_y N_y}{\nu_s A_\delta} \tag{29}$$

(iv) Spillover ratio $|\beta_\delta^{IV}| / \beta_y^{IV}$ *— market integration test (Table 2, row (iv)):*

$$\frac{|\beta_\delta^{IV}|}{\beta_y^{IV}} = \frac{|\partial \delta^{sc} / \partial \bar{\rho}^{sc}|}{\partial y^{sc} / \partial \bar{\rho}^{sc}} = \frac{\Gamma_\times}{N_y} \in (0, 1) \tag{30}$$

[15] Under perfectly elastic supply, the instrument picks up variation from cross-book demand reallocation. If supply elasticity is finite, idiosyncratic household shocks $u^D_{f'}$ provide an additional source of variation through the dollar price $P_t^{s,\text{USD}}$. Appendix G shows that the IV coefficient identifies the same structural object — N_y / A_δ — under both mechanisms; only the instrument's statistical power differs. Given that USDT/USD deviates from par by only 0.05% on average, the cross-book reallocation channel is likely the dominant source of identifying variation.

The identified coefficients enable quantitative policy analysis: given a policy-induced demand shift Δg, the IV responses predict equilibrium price changes accounting for the intermediary's endogenous response and cross-market spillovers. Appendix B formalizes the OLS bias that the IV overcomes.

Table 2: Summary of IV-identified structural objects

	Coefficient	Formula	Economic interpretation
(i)	β_y^{IV}	$\frac{N_y}{A_\delta}$	Marginal cost of stablecoin intermediation: how much y must rise to clear one additional unit. Larger when Γ_{sc}, $\Gamma_\times$ are high or demand elasticity α_{sc} is low.
(ii)	β_δ^{IV}	$-\frac{\Gamma_\times}{A_\delta}$	Cross-market externality: CIP widening per unit of stablecoin flow. Zero when $\Gamma_\times = 0$ (segmented markets).
(iii)	β_s^{IV}	$\frac{\eta_\delta \Gamma_\times + \eta_y N_y}{\nu_s A_\delta}$	Total exchange rate impact: combines direct substitution from stablecoins to spot (η_y) and indirect CIP channel (η_δ).
(iv)	$\frac{\lvert\beta_\delta^{IV}\rvert}{\beta_y^{IV}}$	$\frac{\Gamma_\times}{N_y} \in (0,1)$	Market integration test: share of stablecoin price pressure that transmits to CIP. Near 0 = segmented; near 1 = integrated.

Note: Each row reports the structural object identified by the corresponding IV coefficient from Proposition 4. All expressions are functions of the intermediation cost parameters (Γ_{sc}, Γ_δ, $\Gamma_\times$), demand elasticities (α_{sc}, $\tilde{\alpha}_\delta$), and substitution parameters (η_y, η_δ, ν_s).

5 Empirical implementation and estimation results

This section takes the identification strategy derived in Section 4 to the data. We first describe the construction of the GIV instrument, then present the first-stage and second-stage results, and finally map the estimated IV coefficients back to the model's structural parameters.

5.1 Implementing GIV

We construct our granular instrument in two steps. First, we isolate idiosyncratic shocks by purging the net inflow rate $g_{f,s,t}$ of common global factors by estimating the following equation for each stablecoin s:

$$g_{f,s,t} = \lambda'_{f,s}\mathbf{F}_t + \text{FE}_{s,t} + \text{FE}_{f,s} + u_{f,s,t} \tag{31}$$

where $\mathbf{F}_t$ represents a vector of common factors (e.g., global risk appetite, broad crypto market trends), $\lambda'_{f,s}$ captures heterogeneous sensitivities to these factors, $\text{FE}_{s,t}$ are stablecoin-time fixed effects to absorb coin-specific common shocks, and $\text{FE}_{f,s}$ are currency-coin fixed effects (partially capturing e.g. regional idiosyncrasies and absorbing level differences across coins).

The residual $u_{f,s,t}$ represents the idiosyncratic component of the net inflow rate for pair (f, s) at time t, orthogonal to global and coin-wide factors.

In a second step, we construct the granular instrument as follows. For a given currency j and stablecoin s, the granular instrument is a size-weighted sum of the idiosyncratic shocks for all other currencies $k \neq j$:

$$Z^{\text{GIV}}_{j,s,t} = \sum_{k \neq j} \omega_{k,s} \cdot u_{k,s,t} \tag{32}$$

The weights $\omega_{k,s}$ are the market shares of currency k for stablecoin s, measured by market capitalization in the sample period. This construction ensures that the instrument is dominated by shocks to large, influential currency pairs, supporting the *relevance* condition. By excluding currency j's own shocks, we ensure that the instrument is uncorrelated with local, idiosyncratic demand shocks in j, supporting the *exclusion* restriction.[16]

We then estimate a two-stage least squares procedure. In the first stage we regress the net inflow rate on the instrument and controls:

$$g_{j,s,t} = \theta \, Z^{\text{GIV}}_{j,s,t} + \kappa' \mathbf{X}_{j,t} + \nu_{j,s,t} \tag{33}$$

where $\mathbf{X}_{j,t}$ includes controls such as the Chicago Board Options Exchange Volatility Index (VIX), dollar index, Bitcoin returns, USDT average price in dollars, estimated factors, and coin-currency fixed effects.

In the second stage we then regress the outcome variable $Y_{j,s,t+h}$ on the predicted values $\hat{g}_{j,s,t}$ in the first stage:

$$Y_{j,s,t+h} = \beta^h \, \hat{g}_{j,s,t} + \gamma \mathbf{X}_{j,s,t} + \phi_{j,s} + \zeta_{j,s,t+h} \tag{34}$$

The key identifying assumption is that the granular instrument $Z^{\text{GIV}}_{j,s,t}$ affects the local outcome $Y_{j,s,t+h}$ only through its impact on the local stablecoin inflow rate $g_{j,s,t}$, conditional on controls. This exclusion restriction is plausible because the instrument is built from idiosyncratic shocks in other currencies, purged of common factors, and should have no direct channel to affect the local FX market except via the equilibrium of the global stablecoin market. We assess the strength of the first stage with standard F-statistics and conduct robustness checks by varying the construction of the instrument, using alternative exclusion restrictions (e.g., leave-one-out by coin or by region), and controlling for factors that could produce endogenous local flow responses.

[16] A potential threat to the exclusion restriction could arise if for example global crypto developments would propagate to FX markets independently of stablecoin relative supply factor, e.g. a Bitcoin crash affecting FX through risk sentiment. While such a scenario is highly unlikely, the purging we do in terms of factors and fixed effects, as well as the controls included (e.g. VIX, dollar index, BTC returns) are a first line defense.

GIV relevance. The results of the first-stage regression are presented in Table 3. Our granular instrument is a strong and highly significant predictor of net stablecoin inflows across all specifications, which vary the number of factors used to construct the instrument (columns 1-3) and the set of controls (columns 4-6). The robust F-statistics are well above conventional thresholds, indicating a strong instrument and mitigating concerns about weak identification.

Table 3: First-stage regressions results

	(1)	(2)	(3)	(4)	(5)	(6)
GIV	-0.153***	-0.111***	-0.067***	-0.152***	-0.113***	-0.070***
	(0.0160)	(0.0151)	(0.0164)	(0.0169)	(0.0160)	(0.0169)
Fiat-specific factor 1	1.054***	1.051***	1.036***	1.049***	1.049***	1.036***
	(0.0063)	(0.0040)	(0.0042)	(0.0061)	(0.0039)	(0.0043)
Observations	42,142	42,142	42,142	40,160	40,160	40,160
R-squared	0.215	0.312	0.358	0.219	0.315	0.362
Number of factors	1	2	3	1	2	3
Additional controls	No	No	No	Yes	Yes	Yes
Robust F	14353.5	24058.2	15647.6	2839.0	6324.9	4759.7
K-P rk Wald F	90.46	53.55	16.71	80.50	49.44	17.23
Cragg-Donald F	164.4	99.2	29.4	153.8	98.3	30.8
Adj. R2	0.214	0.311	0.357	0.218	0.314	0.361

Standard errors clustered by date in parentheses. * p<0.10, ** p<0.05, *** p<0.01

Additional controls include log VIX, log of nominal broad dollar index, monthly bitcoin return measured by 30-day change in price, and USDT price in USD (with 5 days of lags).

Note: This table shows the results of regressing net inflow rates of stablecoins (percentage) on our granular instrumental variable. Robust standard errors in parentheses; ***, ** and * respectively denote statistical significance at the 1%, 5% and 10% level. The baseline pools all four USD-pegged stablecoins (USDT, USDC, DAI, and BUSD).

5.2 Estimation of causal effects

Figure 5 presents the core results. The panels plot the dynamic response of our key outcome variables to a one standard deviation shock to the net stablecoin inflow rate, instrumented by our GIV.

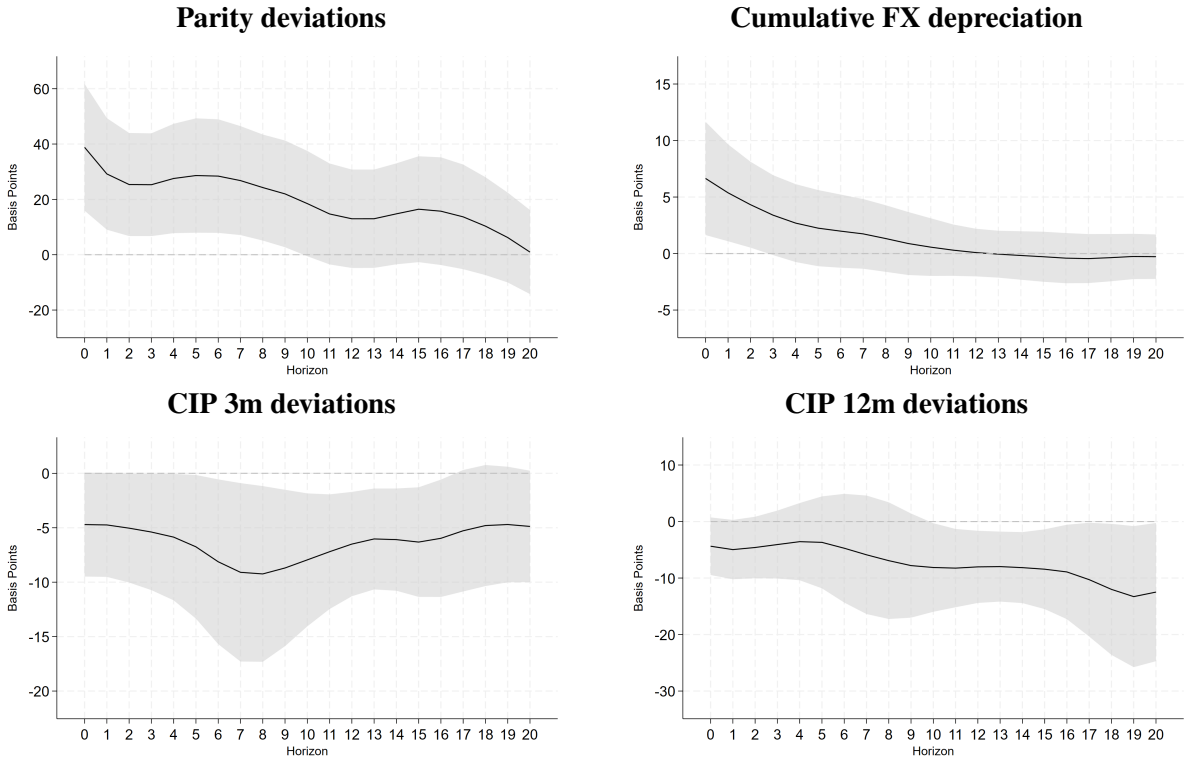
A positive shock to stablecoin inflows causes an immediate, sharp increase in parity deviations of approximately 40 basis points on impact. This effect is highly persistent, decaying only gradually over the following ten days. The persistence of these deviations is consistent with impaired arbitrage of the type that characterizes segmented markets (Gabaix and Maggiori, 2015).

Stablecoin inflows also lead to a statistically significant depreciation of the local currency in

the traditional spot market. The magnitude of the effect declines over time and loses statistical significance after a few days. This confirms that shocks originating in the stablecoin market spill over into traditional FX markets, exerting tangible pressure on the local currency.

Finally, the shock to stablecoin inflows causes a significant decline in 3-month CIP deviations, indicating a deterioration in synthetic funding conditions (a higher dollar premium). In contrast, the effect on 12-month CIP deviations is negligible in the short-term and statistically insignificant. This tenor-specific pattern suggests that arbitrageurs' constraints are more binding at shorter horizons, making short-term CIP deviations more sensitive.[17]

Figure 5: Effect of a one percent increase in instrumented stablecoin net inflows.



The estimated coefficients have direct structural interpretations through Proposition 4. The parity deviation response $\hat{\beta}_y^{IV} \approx 40$ basis points per one percent flow shock measures the marginal cost of stablecoin intermediation: the premium the intermediary requires to absorb an additional unit of stablecoins when facing a supply reduction. The CIP response $\hat{\beta}_\delta^{IV} \approx -17.5$ basis points measures the cross-market externality: the extent to which stablecoin intermediation crowds out synthetic dollar provision through the intermediary's balance sheet. That this coefficient is significantly different from zero constitutes direct evidence against full market segmentation — under $\Gamma_\times = 0$, Proposition 4(ii) predicts no CIP response to stablecoin flow

[17] It is possible that there is heterogeneity across different dimensions in countries' fundamentals. These could include the type of exchange rate regime, the degree of capital account openness or exchange rate restrictions, and the level of financial development. We leave such analysis to further research.

shocks. The exchange rate response $\hat{\beta}_s^{IV} \approx 5 - 10$ basis points captures the total FX impact, combining substitution away from stablecoins toward spot dollar purchases (weighted by η_y) and substitution away from swaps toward spot purchases driven by the CIP widening (weighted by η_δ).

The ratio of the CIP to parity deviation coefficients provides a model-free test of market integration (Proposition 4(iv)). Our estimates yield $|\hat{\beta}_\delta^{IV}|/\hat{\beta}_y^{IV} \approx 0.44$, indicating that roughly 44 percent of the stablecoin parity deviation response transmits to CIP deviations through the balance sheet channel. This value lies well between the polar cases of fully segmented markets (ratio of zero, where stablecoin stress remains contained) and fully integrated markets (ratio approaching one, where stablecoin and swap markets move nearly in lockstep). The intermediate value is consistent with cross-market frictions that are economically significant but do not fully segregate the two markets – a finding we exploit in the counterfactual exercises of Section 6.

The model in Section 4 is static and therefore speaks directly to the effects at impact (i.e., $h = 0$). The persistence of the empirical responses over subsequent days is consistent with two forces that the static framework abstracts from: autocorrelation in stablecoin demand shocks and intermediary balance sheet dynamics that propagate initial shocks over time. We formalize both channels in the dynamic extension of Section 6.3, which generates impulse responses whose shape and half-life closely match the empirical patterns in Figure 5.

6 Counterfactuals

In this section we use the model to conduct counterfactual exercises to quantify how market structure affects the transmission of stablecoin demand shocks to traditional foreign exchange markets. We present the calibration strategy, counterfactuals using the baseline model above, and then the main results of a dynamic extension of the model that generates persistent effects via balance sheet dynamics.[18]

6.1 Calibration

We calibrate the model to match two sets of empirical moments. The first set consists of our GIV estimates of the effects of stablecoin flow shocks on the stablecoin premium, the CIP deviation, and the exchange rate. The second set draws on estimates from the CIP microstructure literature regarding swap market intermediation frictions. We proceed by first fixing the demand-side and pass-through parameters from the literature, and then using the empirical targets to identify the supply-side cost parameters that govern intermediation and cross-market spillovers.

[18] Mapping the effects of the mechanisms we unveil to real macroeconomic data at a monthly or quarterly frequency is left for future research.

Demand-side and pass-through parameters. We set the stablecoin demand elasticity $\alpha_{sc} = 0.55$, the midpoint of own-price elasticities (-0.29 to -0.83) estimated by Benetton and Compiani (2024) for differentiated cryptocurrencies, consistent with the relatively inelastic, necessity-driven demand documented in developing economies (Ahmad et al., 2025). Since α_{sc} does not enter any IV coefficient (Proposition 4), it plays no role in pinning down cost parameters but does affect equilibrium responses through the split between price and quantity adjustment.

The composite swap demand elasticity $\tilde{\alpha}_\delta \equiv \alpha_\delta + \alpha_L = 0.50$ reflects the institutional nature of FX swap demand, where banks and corporates hedge under regulatory mandates that render demand largely insensitive to the prevailing basis (Sushko et al., 2016; Borio et al., 2016), consistent with the sub-unity elasticities in **?**. Although $\tilde{\alpha}_\delta$ enters the IV coefficients through A_δ, its effect is negligible: varying it across $[0.25, 0.75]$ shifts β_δ^{IV} by less than half a basis point.

The swap-to-spot pass-through $\eta_\delta \approx 0.48$ is calibrated from the triangular relationship in Avdjiev et al. (2019), who estimate that a one percentage point broad dollar appreciation is associated with a 2.1 basis point widening of the cross-currency basis; inverting this and normalizing spot market depth $\nu_s = 1$ yields the reported value. This parameter is not precisely identified in the data and should be interpreted as capturing average pass-through, a channel likely stronger in emerging markets with thinner spot liquidity.

Swap market intermediation cost. We calibrate Γ_δ using the microstructure estimates of Rime et al. (2022), who find that a one standard deviation order flow shock widens the cross-currency basis by 0.4–2.7 basis points (normal to stress conditions). Scaling to 1% of daily inter-dealer volume yields a semi-elasticity of 4–27 basis points per percentage point of flow. We target the mid-range of 15.5 basis points for the swap market's own-price response N_δ / A_{sc}.[19] This near-vertical supply schedule is consistent with the quarter-end basis spikes documented by Du et al. (2018).

Supply-side parameters from the GIV identification. The remaining three parameters, the stablecoin-specific intermediation cost Γ_{sc}, the cross-market intermediation cost $\Gamma_\times$, and the stablecoin-to-spot substitution elasticity η_y, are pinned down by the three IV coefficients estimated in Section 5. A 1% instrumented stablecoin flow shock produces a stablecoin parity deviation of 40 basis points, a CIP deviation of -6 basis points, and a spot depreciation of approximately 6 basis points. Together with the structural expressions in Proposition 4, these three moments yield a system of three equations in three unknowns. Intuitively, the CIP spillover coefficient is most informative about the cross-market cost $\Gamma_\times$, as shared balance sheet frictions

[19] The own-price response N_δ / A_{sc} captures the total equilibrium effect of a swap demand shock on the CIP deviation, incorporating both the direct cost of expanding swap intermediation and the general-equilibrium feedback through cross-market linkages.

are the sole channel through which stablecoin shocks transmit to the basis. The stablecoin own-price coefficient then primarily identifies the stablecoin-specific cost Γ_{sc}, which governs how costly it is to intermediate stablecoins per se. Finally, the exchange rate coefficient pins down the stablecoin-to-spot substitution η_y, since the other quantities entering the exchange rate expression are already determined. The resulting $\eta_y = 0.078$ is an order of magnitude smaller than the swap-to-spot substitution η_δ, consistent with stablecoin markets having a much smaller footprint in spot FX turnover than the swap market.

Table 4 reports the full set of parameters. Several features of the calibration merit discussion. The stablecoin-specific intermediation cost $\Gamma_{sc} = 0.337$ is roughly 3.6 times larger than the swap-specific cost $\Gamma_\delta = 0.093$. This ordering is consistent with stablecoin markets involving additional frictions relative to the mature and liquid FX swap infrastructure, including on-ramp and off-ramp costs, blockchain transaction fees, and issuer counterparty risk. The cross-market intermediation cost $\Gamma_\times = 0.065$ is smaller than either own-market cost, implying that intermediaries face larger frictions within each market segment than across them. The ratio $\Gamma_\times/\Gamma_\delta \approx 0.70$ nonetheless indicates that cross-market spillovers represent a meaningful channel through which stablecoin activity affects traditional funding conditions.

Table 4: Model calibration

	Parameter	Value	Source	Target / Rationale
Panel A: Supply-side parameters from GIV identification				
Γ_{sc}	Stablecoin interm. cost	0.337	GIV	$\beta_y^{IV} = 40$ bps
$\Gamma_\times$	Cross-mkt interm. cost	0.065	GIV	$\lvert\beta_\delta^{IV}\rvert = 6$ bps
η_y	SC-to-spot subst.	0.078	GIV	$\beta_s^{IV} \approx 6$ bps
Panel B: Supply-side parameter from CIP literature				
Γ_δ	Swap interm. cost	0.093	RSS22	$N_\delta/A_{sc} = 15.5$ bps
Panel C: Demand-side and pass-through parameters				
α_{sc}	SC demand elasticity	0.55	BKS24	Midpoint of $[0.29, 0.83]$
$\tilde{\alpha}_\delta$	Swap demand elasticity	0.50	GM24	Sub-unity; inelastic demand
η_δ	Swap-to-spot subst.	0.48	AHR19	Triangular relationship
ν_s	Spot market depth	1.00	—	Normalization

Notes: Panel A reports parameters identified from the three GIV coefficients estimated in Section 5. Panel B reports the swap-specific intermediation cost calibrated to match microstructure estimates of the CIP semi-elasticity. Panel C reports demand-side and pass-through parameters fixed from the literature. RSS22: Rime et al. (2022); BKS24: Benetton and Compiani (2024); GM24: **?**; AHR19: Avdjiev et al. (2019). See text for details.

6.2 Counterfactual analysis

In this subsection we examine how price responses to stablecoin demand shocks vary with market structure. For each counterfactual exercise, we vary one parameter while holding all

others at their calibrated values and compute the resulting price responses to a 1% stablecoin flow shock.

Baseline results. Under the calibrated parameters, a 1% stablecoin flow shock, supply response generates a stablecoin parity deviation of $\Delta y = 40$ basis points, a CIP deviation of $\Delta\delta = -6$ basis points (a higher dollar premium), and an exchange rate depreciation of $\Delta s = 6$ basis points. Table 5 reports counterfactual results for variations in the key parameters.

Panel A: Counterfactual equilibrium ($\partial \text{price}/\partial \bar{\rho}^{sc}$, per 1% demand shock)

Table 5: Response to 1% Stablecoin Flow Shock under Different Cross-Market Costs ($\Gamma_\times$)

Scenario	Δy (bps)	$\Delta\delta$ (bps)	Δs (bps)	Interpretation
Baseline	32.8	−4.9	4.9	Calibrated model
$2 \times \Gamma_\times$	36.7	−9.3	7.3	Higher balance sheet linkages
$0.5 \times \Gamma_\times$	30.7	−2.5	3.6	Lower balance sheet linkages
$\Gamma_\times = 0$	28.4	0.0	2.2	Segmented markets

Panel B: Supply response (β^{IV}, per 1% instrumented flow)

Scenario	β_y^{IV} (bps)	β_δ^{IV} (bps)	β_s^{IV} (bps)	$\lvert\beta_\delta\rvert/\beta_y$	Interpretation
Baseline	40.0	−6.0	6.0	0.151	Calibrated model
$2 \times \Gamma_\times$	45.9	−11.7	9.2	0.255	Higher balance sheet linkages
$0.5 \times \Gamma_\times$	36.9	−3.1	4.3	0.083	Lower balance sheet linkages
$\Gamma_\times = 0$	33.7	0.0	2.6	0.000	Segmented markets

Notes: This table reports price responses to a 1% stablecoin flow shock under alternative values of the cross-market cost $\Gamma_\times$, holding all other parameters at their calibrated values. Panel A reports equilibrium price derivatives with respect to the exogenous demand shifter $\bar{\rho}^{sc}$. Panel B reports IV coefficients, to express supply responses per unit of flow g. Δy denotes the stablecoin parity deviation, $\Delta\delta$ denotes the CIP deviation (negative values indicate a higher dollar premium), and Δs denotes the exchange rate depreciation, all in basis points.

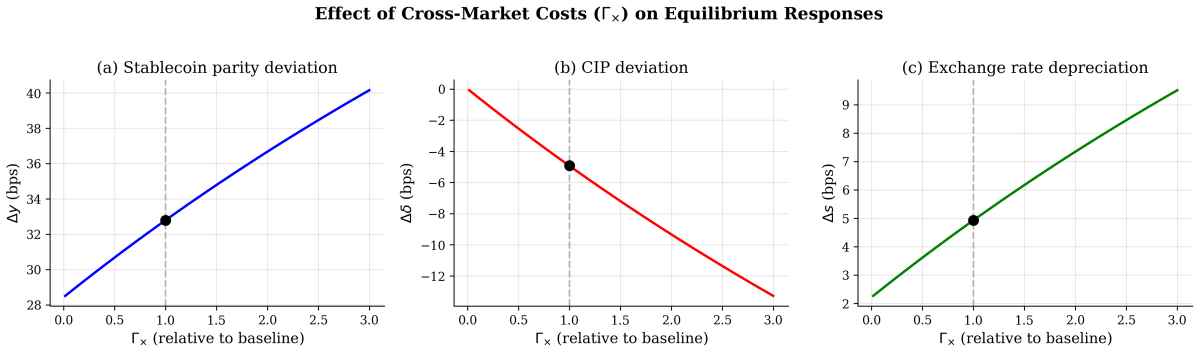
Cross-market costs. Table 5 shows that cross-market costs are the primary driver of spillovers to traditional foreign exchange markets. Doubling $\Gamma_\times$ nearly doubles the CIP spillover and substantially increases all price responses. Concretely, doubling $\Gamma_\times$ corresponds to a scenario in which intermediaries face twice the marginal cost of holding combined stablecoin and swap positions in the same currency — for example, because tighter regulatory capital rules treat stablecoin exposures and FX swap exposures as additive rather than diversifiable, or because a deterioration in counterparty credit conditions raises the cost of netting across desks. In such an environment, every additional dollar of stablecoin intermediation crowds out swap provision more aggressively. As a result, Δy rises from 40 to 46 basis points, $\Delta\delta$ falls from −6 to −12 basis points (a larger dollar premium), and Δs rises from 6 to 9 basis points. In contrast, if

markets were fully segmented ($\Gamma_\times \to 0$), the CIP spillover would vanish completely, while the exchange rate spillover would half.

The intuition follows directly from the model structure. When $\Gamma_\times > 0$, intermediaries face a cost that depends on their total currency exposure across both stablecoin and swap markets. A stablecoin demand shock that expands intermediaries' stablecoin positions therefore raises the marginal cost of swap provision, reducing swap supply and lowering CIP deviations (increasing the dollar premium). The larger is $\Gamma_\times$, the stronger this balance sheet channel becomes. In the limiting case where $\Gamma_\times = 0$, the two markets are effectively segmented, and stablecoin shocks affect stablecoin prices but do not transmit to swap markets.

Figure 6 illustrates these effects. The stablecoin parity deviation and exchange rate depreciation increase monotonically in $\Gamma_\times$, while the CIP deviation becomes more negative (a larger dollar premium). The CIP deviation exhibits the steepest response. This pattern confirms that spillovers to traditional foreign exchange markets operate primarily through the cross-market cost channel.

Figure 6: Effect of cross-market costs ($\Gamma_\times$) on Price Responses.



Notes: This figure displays equilibrium price responses to a 1% stablecoin flow shock across different cross-market cost parameter $\Gamma_\times$. The left panel shows the stablecoin parity deviation (Δy), the middle panel shows the CIP deviation ($\Delta\delta$, where more negative values indicate a higher dollar premium), and the right panel shows the exchange rate depreciation (Δs). The horizontal axis measures $\Gamma_\times$ relative to its calibrated baseline value. Dots indicate baseline values.

The cross-market cost parameter $\Gamma_\times$ also provides a natural interpretation for cross-country heterogeneity. The cross-country heterogeneity documented in our stylized facts can be rationalized through variation in effective $\Gamma_\times$ across countries. In economies with open capital accounts, global intermediaries can offset locally-acquired currency exposure through their international operations. This ability to manage currency risk globally reduces the effective cross-market cost. By contrast, in economies with binding capital controls, intermediaries face regulatory and legal barriers to transferring positions offshore. Capital controls therefore raise the effective $\Gamma_\times$ by preventing the international diversification of intermediary balance sheets. The counterfactual results in Panel A indicate that such variation in $\Gamma_\times$ can generate substantial differences in spillover magnitudes: moving from a low-$\Gamma_\times$ environment (open capital account)

to a high-$\Gamma_\times$ environment (binding capital controls) could more than double the CIP spillover from stablecoin shocks.

Appendix E discusses the effects of increasing stablecoin-specific costs (Γ_{sc}) and of increasing stablecoin demand elasticity (α_{sc}).

Policy implications. The counterfactual analysis yields two policy implications. First, the magnitude of spillovers depends critically on intermediary balance sheet linkages. Policies that affect $\Gamma_\times$, such as capital requirements that treat stablecoin and swap exposures as substitutes, will directly influence how stablecoin market stress transmits to traditional foreign exchange markets. Specifically, capital frameworks that permit netting of stablecoin and FX swap exposures would lower the effective $\Gamma_\times$ and dampen spillovers, while ring-fencing crypto activities into separately capitalized entities would eliminate the balance sheet link altogether (at the cost of reducing cross-market liquidity provision). Second, the finding that CIP spillovers nearly vanish when $\Gamma_\times \to 0$ suggests that market segmentation, while potentially costly in terms of allocative efficiency, would limit contagion from stablecoin markets to traditional FX markets.

6.3 Dynamic counterfactuals

Dynamic model. This subsection extends the baseline model to incorporate demand persistence and intermediary balance sheet dynamics,[20] delivering three results. First, the dynamic model generates impulse responses whose shape and persistence closely match our empirical estimates. Second, it shows that spillovers are state-dependent: identical shocks produce larger effects when intermediary balance sheets are impaired. Third, doubling cross-market frictions amplifies spillovers more than proportionally through a feedback loop between prices and wealth. Appendix F provides the full specification.

There are two new assumptions. First, stablecoin demand $\bar{D}_t^{sc}$ follow an AR(1) process that generates autocorrelation so shocks decay gradually rather than revert immediately. Second, wealth (W_t) is a key state variable for the intermediary's problem: she evaluates her portfolio composition in synthetic dollars and stablecoins relative to capital and constrained by the law of motion of wealth.

Intermediary problem with wealth-dependent risk capacity. The representative intermediary enters period t with wealth W_t and chooses positions in synthetic dollar intermediation

[20]The extension draws on the intermediary asset pricing literature, particularly He and Krishnamurthy (2013), who demonstrate that when financial intermediaries face equity capital constraints, risk premia become state-dependent and exhibit nonlinear dynamics during crises. In the international finance context, Gabaix and Maggiori (2015) show that capital flows drive exchange rates by altering the balance sheets of financiers who bear currency risk, while Du et al. (2018) document that CIP deviations reflect the shadow cost of intermediary balance sheet capacity. Our dynamic extension captures these mechanisms in a tractable framework that generates persistent spillovers and state-dependent amplification.

(B_t^{syn}) and stablecoin intermediation (B_t^{sc}). The key departure from the static model is that the intermediary's costs are scaled by wealth, capturing the idea that positions must be evaluated relative to capital. The intermediary solves:

$$\max_{B_t^{syn}, B_t^{sc}} \quad \delta_t B_t^{syn} + y_t B_t^{sc} - \frac{\Gamma_\delta}{2W_t}(B_t^{syn})^2 - \frac{\Gamma_{sc}}{2W_t}(B_t^{sc})^2 - \frac{\Gamma_\times}{2W_t}(B_t^{syn} + B_t^{sc})^2 \tag{35}$$

The scaling by W_t in the cost terms admits a natural interpretation: a given dollar position represents greater risk exposure when measured against a smaller capital base. An intermediary with $10 billion in capital can comfortably absorb losses on a $1 billion position; the same position would be precarious for an intermediary with only $2 billion. This formulation is analogous to measuring portfolio risk relative to equity, as in the Value-at-Risk (VaR)-constraint literature (Brunnermeier and Pedersen, 2009; Adrian and Shin, 2014), and captures the risk-bearing capacity mechanism emphasized by Gabaix and Maggiori (2015).

Wealth dynamics. Intermediary wealth evolves with portfolio returns:

$$W_{t+1} = W_t \left(1 + \frac{B_t^{syn}}{W_t}\delta_t + \frac{B_t^{sc}}{W_t} y_t - \phi \cdot \Delta s_t\right) = W_t \left(1 + b_t^{syn}\delta_t + b_t^{sc} y_t - \phi \cdot \Delta s_t\right) \tag{36}$$

where ϕ captures the intermediary's net foreign exchange exposure and Δs_t is the rate of exchange rate depreciation. The position-to-wealth ratios b_t^{syn} and b_t^{sc} determine how portfolio returns translate into wealth growth.

This specification creates scope for amplification. Adverse shocks that raise the stablecoin parity deviation y_t and lower δ_t (increasing the dollar premium) may generate mark-to-market losses that deplete wealth. With reduced wealth, the intermediary's risk-bearing capacity falls, positions contract, and prices must adjust further to clear markets. This feedback loop is the central mechanism in He and Krishnamurthy (2013): losses reduce capacity, and reduced capacity sustains elevated prices.

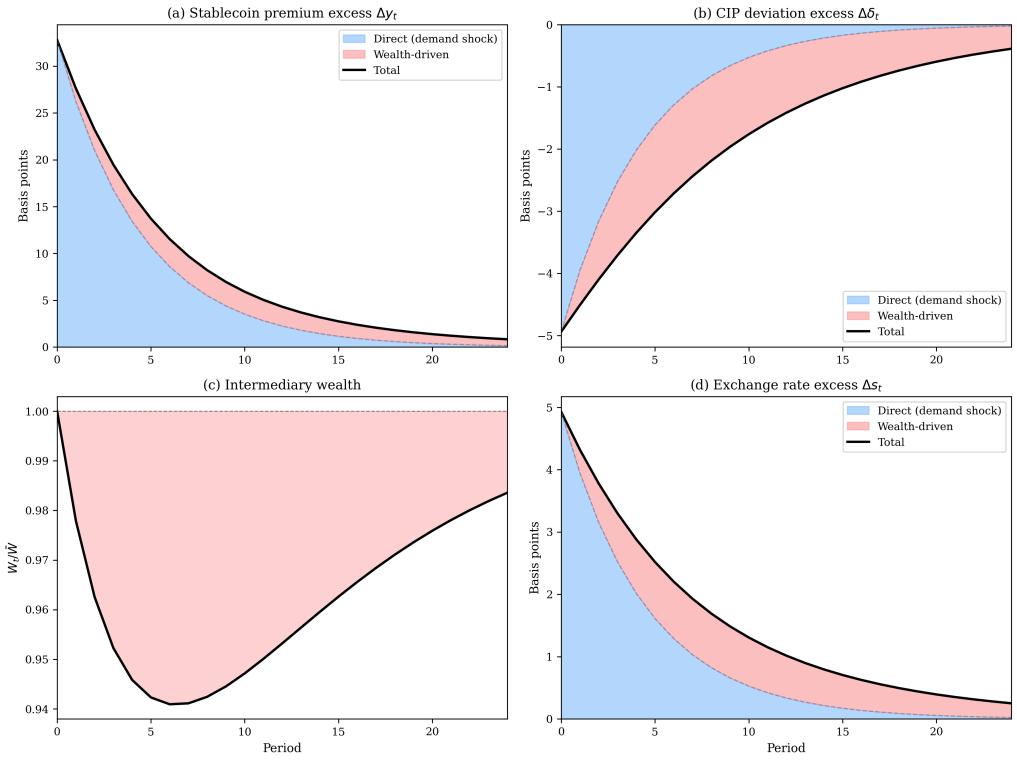
In equilibrium, prices (y_t, δ_t, s_t) are a function of W_t and stablecoin demand $\bar{D}_t^{sc}$ given the cost parameters Γ. The key comparative statics are $\partial y_t / \partial W_t < 0$ and $\partial \delta_t / \partial W_t > 0$: lower intermediary wealth implies higher stablecoin parity deviations and lower CIP deviations (a higher dollar premium). Intuitively, when the intermediary has less capital, it supplies less risk-bearing capacity to the market. With reduced supply of intermediation, prices must adjust to equilibrate demand.

Figure 7 displays the impulse response to a one percent stablecoin demand shock relative to the initial steady state.[21] The shock generates impact effects of 34 basis points for the stablecoin

[21] Appendix F.3 presents the calibration. Static parameters are taken from the main text, while dynamic parameters (stablecoin demand persistence, FX exposure, and steady-state wealth) are chosen to match empirical features of stablecoin markets and intermediary balance sheets.

parity deviation y, -5 basis points for the CIP deviation δ (a higher dollar premium), and 5 basis points for exchange rate depreciation s. These impact effects equal those of the static model, since at steady-state wealth the dynamic model nests the static framework.

Figure 7: Impulse response to a one percent stablecoin demand shock from steady state.

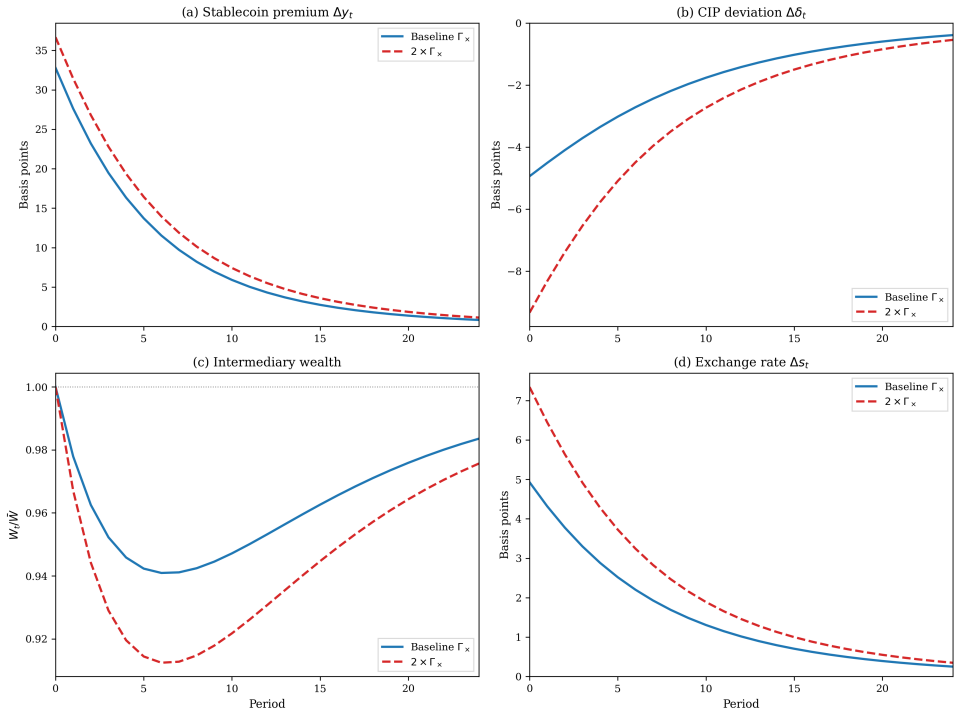


Note: The upper panels display the stablecoin parity deviation (left) and CIP deviation (right, where negative values indicate a higher dollar premium) in basis points. The lower panels show exchange rate depreciation (left, basis points) and intermediary wealth (right, percent deviation from steady state). The blue shade represents the direct effect of stablecoin demand shocks, and the pink shade shows the amplification through intermediaries' wealth effects.

A key prediction of the model is that identical shocks generate larger price responses when intermediary balance sheets are impaired. This property follows directly from the equilibrium structure: with the stablecoin parity deviation scaling as $1/W_t$ and the CIP deviation scaling as W_t, lower wealth translates mechanically into higher stablecoin parity deviations, larger dollar premia, and further FX depreciation. Endogenous wealth depletion introduce further persistence of demand shocks and its spillovers, in line with our empirical findings.

Dynamic response to market friction shocks. The static analysis identified the cross-market cost $\Gamma_\times$ as the key parameter governing spillovers from stablecoin markets to CIP deviations. The dynamic model amplifies this role through the balance sheet channel. Figure 8 compares impulse responses under the baseline calibration to a counterfactual with doubled cross-market costs.

Figure 8: Impulse responses under baseline and doubled cross-market costs ($2 \times \Gamma_\times$)



Notes: Solid lines show baseline responses; dashed red lines show responses under higher cross-market frictions.

Doubling $\Gamma_\times$ increases the impact effect on the stablecoin parity deviation by about 5 basis points. The CIP spillover increases from -6 to -9 basis points, nearly a 50 percent increase in the dollar premium. The more-than-proportional increase in CIP spillovers reflects compounding through the balance sheet channel: higher $\Gamma_\times$ generates larger initial price responses, which cause greater wealth depletion, which further amplifies subsequent price dynamics.

The dynamic model strengthens the case for policy attention to cross-market frictions. The dynamic extension shows that reducing these frictions yields compounding benefits: smaller spillovers preserve intermediary capital, which sustains risk-bearing capacity, which further dampens price responses to future shocks.

Run dynamics and amplification. In an additional extension (Appendix F.8), we incorporate redemption frictions that amplify during stress. When flow shocks coincide with friction shocks (as during the Terra/Luna collapse or the SVB-induced USDC depeg) the interaction between redemption costs and balance sheet depletion can multiply baseline spillover effects severalfold, underscoring that stablecoin-to-FX transmission can escalate sharply during crisis episodes.

7 Conclusion

This paper provides new systematic evidence of a causal link between the crypto-based financial system and traditional foreign exchange markets. Using granular data on four USD-pegged stablecoins and 27 fiat currencies, we document substantial parity deviations between stablecoin and traditional FX venues and show that exogenous shocks to stablecoin inflows, identified via a granular instrumental variable strategy, cause these deviations to widen, depreciate the local currency, and impair synthetic dollar funding conditions. A model of segmented arbitrage rationalizes these findings and provides structural foundations for the identification strategy: the IV coefficients map directly to the marginal cost of intermediation, the cross-market spillover intensity, and a market integration test.

The spillover ratio we estimate (approximately 0.15) indicates that roughly one-sixth of stablecoin price pressure transmits to synthetic dollar funding costs, affecting banks, corporates, and sovereigns that rely on FX swaps regardless of whether they participate in crypto markets. Counterfactual exercises show that cross-market frictions are the primary determinant of spillover magnitude, and a dynamic extension reveals that propagation is state-dependent: depleted intermediary capital amplifies price responses, and when flow shocks coincide with redemption frictions, spillover effects can reach several multiples of the baseline.

These results point to two policy directions. First, prudential requirements on stablecoin intermediaries, such as capital buffers, reserve liquidity mandates, and limits on concentrated currency exposures, would shrink the spillover channel at its source by reducing cross-market frictions and preserving risk-bearing capacity. Second, policymakers responsible for currency stability, particularly in emerging markets, could consider incorporating stablecoin market monitoring into their macroprudential surveillance, since the transmission mechanism runs through balance sheets that also serve traditional FX swap markets. Monitoring the evolution of the spillover ratio as stablecoin markets grow would provide an early warning indicator of increasing interconnectedness.

Our work opens several avenues for future research, including investigating how spillover magnitudes vary with capital account openness, exchange rate regime, and local FX market depth; analyzing how stablecoin design features (algorithmic vs. reserve-backed vs. CBDC) map into different intermediation costs and spillover profiles; tracking the spillover ratio over time as a gauge of crypto–FX integration; and studying whether stablecoin flows weaken the effectiveness of monetary policy and capital flow management measures in economies where they offer an alternative channel for cross-border capital movements.

References

Adrian, T. and H. S. Shin (2014): "Procyclical Leverage and Value-at-Risk," *Review of Financial Studies*, 27, 373–403.

Ahmad, W., A. M. Kutan, and S. Gupta (2025): "Bitcoin Adoption and Price Elasticity of Demand: Cross-Country Insights," *Humanities and Social Sciences Communications*, 12, Article 1028.

Ahmed, R. and I. Aldasoro (2025): "Stablecoins and safe asset prices," BIS Working Paper 1279.

Ahmed, R., I. Aldasoro, and C. Duley (2025): "Public information and stablecoin runs," BIS Working Paper 1164.

Auer, R., U. Lewrick, and J. Paulick (2025): "DeFiying gravity? An empirical analysis of cross-border Bitcoin, Ether and stablecoin flows," BIS Working Paper 1265.

Avdjiev, S., W. Du, C. Koch, and H. S. Shin (2019): "The dollar, bank leverage, and deviations from covered interest parity," *American Economic Review: Insights*, 1, 193–208.

Azar, P. D., G. Baughman, F. Carapella, J. Gerszten, A. Lubis, J. P. Perez-Sangimino, D. E. Rappoport, C. Scotti, N. Swem, A. Vardoulakis, et al. (2022): "The financial stability implications of digital assets," Tech. rep., Staff Reports.

Barthelemy, J., P. Gardin, and B. Nguyen (2023): "Stablecoins and the Financing of the Real Economy," Banque de France Working Paper 908.

Baughman, G., F. Carapella, J. Gerszten, and D. C. Mills (2022): "The Stable in Stablecoins," FEDS Notes 2022-12-17, Board of Governors of the Federal Reserve System.

Benetton, M. and G. Compiani (2024): "Investors' Beliefs and Cryptocurrency Prices," *The Review of Asset Pricing Studies*, 14, 197–236.

Bertsch, C. (2023): "Stablecoins: Adoption and Fragility," Working Paper Series 423, Sveriges Riksbank (Central Bank of Sweden).

Borio, C., R. McCauley, P. McGuire, and V. Sushko (2016): "Covered interest parity lost: understanding the cross-currency basis," *BIS Quarterly Review*, 45–64.

Brunnermeier, M. K. and L. H. Pedersen (2009): "Market Liquidity and Funding Liquidity," *Review of Financial Studies*, 22, 2201–2238.

D'Avernas, A., V. Maurin, and Q. Vandeweyer (2023): "Can Stablecoins Be Stable?" Working paper, (available at SSRN).

Du, W., A. Tepper, and A. Verdelhan (2018): "Deviations from Covered Interest Rate Parity," *The Journal of Finance*, 73, 915–957.

Franz, F.-C. and A. Valentin (2020): "Crypto covered interest parity deviations," *Available at SSRN 3702212*.

Gabaix, X. and R. S. J. Koijen (2024): "Granular Instrumental Variables," *Journal of Political Economy*, 132, 2274–2303.

Gabaix, X. and M. Maggiori (2015): "International Liquidity and Exchange Rate Dynamics," *The Quarterly Journal of Economics*, 130, 1369–1420.

Goel, T., U. Lewrick, and I. Agarwal (2025): "Making stablecoins stable(r): Can regulation help?" *Available at SSRN: https://ssrn.com/abstract=5070116*.

Gorton, G. B., E. C. Klee, C. P. Ross, S. Y. Ross, and A. P. Vardoulakis (2026): "Leverage and Stablecoin Pegs," *Journal of Financial and Quantitative Analysis*, 61, 99–136.

Gorton, G. B., C. P. Ross, and S. Y. Ross (2025): "Making Money," Working paper.

Graf von Luckner, C., C. M. Reinhart, and K. Rogoff (2023): "Decrypting new age international capital flows," *Journal of Monetary Economics*, 138, 104–122.

Graf von Luckner, C. M., R. Koepke, and S. Sgherri (2024): "Crypto as a marketplace for capital flight," IMF Working Papers 2024/133, International Monetary Fund.

He, Z. and A. Krishnamurthy (2013): "Intermediary Asset Pricing," *American Economic Review*, 103, 732–770.

Itskhoki, O. and D. Mukhin (2021): "Exchange Rate Disconnect in General Equilibrium," *Journal of Political Economy*, 129, 2183–2232.

Kaiko Research (2025): "Stablecoin Dominance & Market Liquidity," .

Kalemli-Özcan, Ş. and L. Varela (2024): "Five Facts about the UIP Premium," *Mimeo*.

Kim, S. (2025a): "How the Cryptocurrency Market is Connected to the Financial Market," *working paper*.

——— (2025b): "Macro-Financial Impact of Stablecoin's Demand for Treasuries," *working paper*.

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

Ma, Y., Z. Yeng, and A. L. Zhang (2023): "Stablecoin runs and the centralization of arbitrage," Working paper, (available at SSRN).

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

Reuter, M. (2025): "Decrypting Crypto: How to Estimate International Stablecoin Flows," IMF Working Paper 2025/141, International Monetary Fund.

Rime, D., A. Schrimpf, and O. Syrstad (2022): "Covered Interest Parity Arbitrage," *The Review of Financial Studies*, 35, 5185–5227.

Schmeling, M., A. Schrimpf, and K. Todorov (2023): "Crypto carry," BIS Working Papers 1087, Bank for International Settlements.

Shleifer, A. and R. W. Vishny (1997): "The Limits of Arbitrage," *The Journal of Finance*, 52, 35–55.

Sushko, V., C. Borio, R. N. McCauley, and P. McGuire (2016): "The failure of covered interest parity: FX hedging demand and costly balance sheets," BIS Working Papers 590, Bank for International Settlements.

A Shock structure and endogeneity

To understand why OLS regressions fail to identify causal effects, and why our instrument succeeds, we must be explicit about the sources of variation driving demand on each stablecoin exchange book. Recall that net demand pressure on book f combines local household demand and cross-book participant demand,

$$\bar{\rho}_f^{sc} = \bar{D}_f^{sc} + S_f^{sc}. \tag{37}$$

We decompose this into common and idiosyncratic components. Local household demand has the structure

$$\bar{D}_f^{sc} = \bar{D}_0 + \phi\mu + u_f^D, \tag{38}$$

where $\bar{D}_0$ is baseline demand common to all countries, $\phi\mu$ captures exposure to global shocks μ that affect all countries simultaneously (such as a flight to dollar safety during global risk-off episodes or a surge in crypto adoption following a major platform launch), and u_f^D is an idiosyncratic local household shock specific to country f, reflecting local currency instability, country-specific regulatory changes, or localized growth in crypto payment infrastructure.

Cross-book participant demand on book f has the decomposition

$$S_f^{sc} = \bar{S}_f^{sc} + v_f, \qquad \sum_f v_f = 0, \tag{39}$$

where $\bar{S}_f^{sc}$ is the baseline allocation and v_f captures idiosyncratic reallocation driven by urgency shocks in other countries. The zero-sum property follows from the budget constraint $\sum_f S_f^{sc} = W$.

Combining these expressions, total demand on book f is

$$\tilde{D}_f^{sc} \equiv \bar{D}_f^{sc} + S_f^{sc} = (\bar{D}_0 + \bar{S}_f^{sc}) + \phi\mu + \tilde{u}_f, \tag{40}$$

where $\tilde{u}_f \equiv u_f^D + v_f$ is the idiosyncratic component of total demand on book f, combining local household shocks and cross-book reallocation.

The identification challenge is that stablecoins are not the only channel through which these shocks affect prices. Global and local shocks also influence stablecoin parity deviations, CIP

deviations, and exchange rates through channels unrelated to stablecoin intermediation,

$$y_f = \underbrace{y_f^{sc}(\bar{\rho}_f^{sc})}_{\text{stablecoin channel}} + \underbrace{\xi_y \mu}_{\text{global shock}} + \underbrace{\varepsilon_f^y}_{\text{local factors}} \tag{41}$$

$$\delta_f = \underbrace{\delta_f^{sc}(\bar{\rho}_f^{sc})}_{\text{stablecoin channel}} - \underbrace{\xi_\delta \mu}_{\text{global shock}} - \underbrace{\varepsilon_f^\delta}_{\text{local factors}} \tag{42}$$

$$s_f = \underbrace{s_f^{sc}(\bar{\rho}_f^{sc})}_{\text{stablecoin channel}} + \underbrace{\xi_s \mu}_{\text{global shock}} + \underbrace{\varepsilon_f^s}_{\text{local factors}} \tag{43}$$

The first term in each equation, derived from the equilibrium in Proposition 1, captures how stablecoin market conditions affect prices through the intermediary's balance sheet. The second term captures direct effects of global shocks. A flight to safety ($\mu > 0$) may independently widen CIP deviations (make δ more negative, hence the negative sign with $\xi_\delta > 0$) as global banks reduce dollar lending, and depreciate emerging market currencies ($\xi_s > 0$) through portfolio rebalancing, even holding fixed the stablecoin channel. The third term captures local factors unrelated to stablecoins, such as domestic monetary policy surprises, local banking sector stress, or terms-of-trade shocks.

We formalize the correlation structure of these shocks.

Assumption 1 (Shock structure)**.** *The shocks satisfy:*

(i) Idiosyncratic total demand shocks are independent across countries: $Cov(\tilde{u}_f, \tilde{u}_{f'}) = 0$ for $f \neq f'$.

(ii) Idiosyncratic total demand shocks are correlated with local unobservables: $Cov(\tilde{u}_f, \varepsilon_f^y) = \sigma_{\tilde{u}\varepsilon}^y \neq 0$, and similarly for δ and s.

(iii) Idiosyncratic total demand shocks in country f are uncorrelated with unobservables in other countries: $Cov(\tilde{u}_f, \varepsilon_{f'}^y) = 0$ for $f \neq f'$.

Condition (i) states that idiosyncratic events driving total demand on one country's book are unrelated to those on another. This follows from two properties. Local household shocks u_f^D are idiosyncratic by nature. Cross-book reallocation shocks v_f are driven by urgency in other countries ($v_f = v_f(\theta_{-f})$), so that v_f and $v_{f'}$ are uncorrelated conditional on having purged the common factor μ.[22] Condition (ii) is the source of endogeneity that invalidates OLS. The same local events that drive stablecoin demand also affect prices through other channels. For example, an episode of local currency instability may simultaneously trigger capital flight into stablecoins (raising $\tilde{u}_f$) and widen CIP deviations through reduced intermediation capacity (contributing to ε_f^δ). Condition (iii) states that while local shocks create correlation between

[22] Formally, $\text{Cov}(\tilde{u}_f, \tilde{u}_{f'}) = \text{Cov}(u_f^D, u_{f'}^D) + \text{Cov}(u_f^D, v_{f'}) + \text{Cov}(v_f, u_{f'}^D) + \text{Cov}(v_f, v_{f'})$. The first three terms are zero by independence of idiosyncratic local shocks and reallocation shocks. The last term is zero after factor purging removes common components, since v_f depends on urgency shocks θ_{-f} while $v_{f'}$ depends on $\theta_{-f'}$.

demand and prices *within* a country, they do not create correlation *across* countries. This restriction enables our identification strategy.

B OLS bias

Consider regressing prices on stablecoin flows $g_f \equiv B_f^{sc}$. To see why OLS is biased, note that flows depend on both demand conditions and equilibrium prices. From market clearing,

$$g_f = \bar{\rho}_f^{sc} - \alpha_{sc} y_f. \tag{44}$$

Substituting the price equation (41) and using $\partial y^{sc}/\partial \bar{\rho}^{sc} = N_y/\Delta$ from Proposition 1,

$$g_f = \gamma_g \bar{\rho}_f^{sc} - \alpha_{sc}(\xi_y \mu + \varepsilon_f^y), \tag{45}$$

where $\gamma_g \equiv 1 - \alpha_{sc}(N_y/\Delta) = A_\delta/\Delta > 0$ captures how much of a demand shock translates into flows after accounting for the endogenous price response. Flows thus inherit variation from three sources: idiosyncratic demand shocks $\tilde{u}_f$ (embedded in $\bar{\rho}_f^{sc}$), global shocks μ, and local price determinants ε_f^y. The latter two create bias.

Proposition 5 (OLS bias)**.** *Under Assumption 1, the OLS coefficient from regressing δ_f on g_f is*

$$\beta_\delta^{OLS} = \underbrace{-\frac{\Gamma_\times}{A_\delta}}_{\beta_\delta^{IV}} + \underbrace{\frac{-\phi \xi_\delta \sigma_\mu^2 - \sigma_{\tilde{u}\varepsilon}^\delta}{Var(g_f)/\gamma_g}}_{bias} \tag{46}$$

where $\sigma_\mu^2 = Var(\mu)$, $\sigma_{\tilde{u}\varepsilon}^\delta = Cov(\tilde{u}_f, \varepsilon_f^\delta)$, and $\gamma_g = A_\delta/\Delta$. Similarly for the stablecoin parity deviation and exchange rate with analogous bias terms.

Proof. See Appendix C. □

The OLS bias can be decomposed into two distinct sources of confounding.

The first is global shock bias, arising from the term $-\phi \xi_\delta \sigma_\mu^2$. Global shocks affect both stablecoin demand and prices through parallel channels. Consider a deterioration in global risk sentiment ($\mu > 0$) that drives investors worldwide into dollar-denominated safe assets including stablecoins ($\phi > 0$). The same risk-off environment independently widens CIP deviations (makes δ more negative, with $\xi_\delta > 0$) as global financial institutions tighten dollar lending and reduce balance sheet exposure to emerging markets. OLS attributes the entire correlation between flows and CIP to the stablecoin channel, overstating the causal effect.

The second is local confounding bias, arising from the term $-\sigma_{\tilde{u}\varepsilon}^\delta$. Local factors that drive total demand on a country's book are often correlated with local price determinants. This

correlation operates through two sub-channels. Through household demand (u_f^D component of $\tilde{u}_f$), a shock that erodes confidence in the domestic financial system causes households to increase stablecoin holdings (u_f^D rises) while the same shock widens CIP deviations as local institutions lose access to dollar funding (ε_f^δ rises, making δ more negative). Through cross-book reallocation (v_f component of $\tilde{u}_f$), an urgency shock in country f draws cross-book participants toward book f (v_f rises) while the same local stress widens CIP. In both cases, OLS conflates the causal effect of stablecoins with correlated responses to the same underlying shock.

The direction of OLS bias depends on the sign and magnitude of these correlations. In crisis episodes, which may generate much of the high-frequency variation in our data, one would typically expect stablecoin demand to rise with global stress as investors seek dollar safety ($\phi > 0$), global stress to independently widen CIP as banks reduce dollar intermediation ($\xi_\delta > 0$), and local stress to simultaneously raise total book demand and widen CIP ($\sigma_{\tilde{u}\varepsilon}^\delta > 0$). This configuration implies that all bias terms push the OLS estimate to be more negative than the true causal effect, so OLS likely *overstates* the spillover from stablecoins to CIP. As such, the estimated effect partially reflects correlated responses to common shocks rather than causal transmission through the stablecoin channel.

However, the bias could work in the opposite direction in other contexts. If stablecoin adoption is driven by fintech innovation ($u_f^D > 0$) in countries experiencing financial deepening and improved market access ($\varepsilon_f^\delta < 0$, reflecting narrowing CIP), then $\sigma_{\tilde{u}\varepsilon}^\delta < 0$ and OLS would *understate* the causal effect. The sign of the bias is ultimately an empirical question, which motivates our instrumental variable strategy that eliminates bias regardless of its direction.

B.1 Proof of Proposition 5

From the price equation (42) and the flow expression,

$$\delta_f = -\frac{\Gamma_\times}{\Delta}\bar{\rho}_f^{sc} + (\text{terms in } \bar{\rho}^{syn}) - \xi_\delta \mu - \varepsilon_f^\delta$$

$$g_f = \gamma_g \bar{\rho}_f^{sc} - \alpha_{sc}(\xi_y \mu + \varepsilon_f^y)$$

where $\bar{\rho}_f^{sc} = (\bar{D}_0 + \bar{S}_f^{sc}) + \phi\mu + \tilde{u}_f$ from the demand decomposition (40). We treat $\bar{\rho}^{syn}$ and $\bar{D}_0 + \bar{S}_f^{sc}$ as non-stochastic conditional on country f's baseline characteristics. The stochastic components of interest are then

$$\delta_f - E[\delta_f] = -\frac{\Gamma_\times}{\Delta}(\phi\mu + \tilde{u}_f) - \xi_\delta \mu - \varepsilon_f^\delta$$

$$g_f - E[g_f] = \gamma_g(\phi\mu + \tilde{u}_f) - \alpha_{sc}(\xi_y \mu + \varepsilon_f^y).$$

The OLS coefficient is $\beta_\delta^{OLS} = \text{Cov}(\delta_f, g_f)/\text{Var}(g_f)$. Computing the numerator,

$$\begin{aligned}\text{Cov}(\delta_f, g_f) = -\frac{\Gamma_\times}{\Delta}\gamma_g(\phi^2\sigma_\mu^2 + \sigma_{\tilde{u}}^2) + \frac{\alpha_{sc}\Gamma_\times}{\Delta}(\phi\xi_y\sigma_\mu^2 + \sigma_{\tilde{u}\varepsilon}^y) \\ - \xi_\delta\gamma_g\phi\sigma_\mu^2 + \alpha_{sc}\xi_\delta\xi_y\sigma_\mu^2 - \gamma_g\sigma_{\tilde{u}\varepsilon}^\delta + \alpha_{sc}\text{Cov}(\varepsilon_f^\delta, \varepsilon_f^y),\end{aligned}$$

where $\sigma_{\tilde{u}}^2 = \text{Var}(\tilde{u}_f)$, $\sigma_{\tilde{u}\varepsilon}^y = \text{Cov}(\tilde{u}_f, \varepsilon_f^y)$, and $\sigma_{\tilde{u}\varepsilon}^\delta = \text{Cov}(\tilde{u}_f, \varepsilon_f^\delta)$.

The variance of flows is

$$\text{Var}(g_f) = \gamma_g^2(\phi^2\sigma_\mu^2 + \sigma_{\tilde{u}}^2) + \alpha_{sc}^2(\xi_y^2\sigma_\mu^2 + \sigma_{\varepsilon^y}^2) - 2\alpha_{sc}\gamma_g(\phi\xi_y\sigma_\mu^2 + \sigma_{\tilde{u}\varepsilon}^y).$$

To separate the causal component from the bias, write

$$\text{Cov}(\delta_f, g_f) = -\frac{\Gamma_\times}{\Delta} \cdot \frac{\text{Var}(g_f)}{\gamma_g} + \left[-\phi\xi_\delta\sigma_\mu^2 - \sigma_{\tilde{u}\varepsilon}^\delta + \text{h.o.t.}\right]$$

where the first term uses the identity $\text{Var}(g_f)/\gamma_g = \gamma_g(\phi^2\sigma_\mu^2 + \sigma_{\tilde{u}}^2) - \alpha_{sc}(\phi\xi_y\sigma_\mu^2 + \sigma_{\tilde{u}\varepsilon}^y) + (\alpha_{sc}^2/\gamma_g)(\xi_y^2\sigma_\mu^2 + \sigma_{\varepsilon^y}^2)$ and the higher-order terms (involving products of global and local parameters) are absorbed into $\text{Var}(g_f)$ upon division. Dividing through,

$$\beta_\delta^{OLS} = \frac{-\Gamma_\times/\Delta}{\gamma_g} + \frac{-\phi\xi_\delta\sigma_\mu^2 - \sigma_{\tilde{u}\varepsilon}^\delta}{\text{Var}(g_f)/\gamma_g}.$$

Using $\gamma_g = A_\delta/\Delta$, the first term simplifies to $-\Gamma_\times/A_\delta = \beta_\delta^{IV}$, yielding the stated expression. □

Remark. The OLS bias for the parity deviation regression follows analogously. The coefficient from regressing y_f on g_f is

$$\beta_y^{OLS} = \underbrace{\frac{N_y}{A_\delta}}_{\beta_y^{IV}} + \frac{\phi\xi_y\sigma_\mu^2 + \sigma_{\tilde{u}\varepsilon}^y}{\text{Var}(g_f)/\gamma_g}. \tag{47}$$

The sign of the global bias differs from the CIP case. Global risk-off shocks ($\phi > 0$, $\xi_y > 0$) push the OLS estimate *above* the true causal effect for the parity deviation, since both stablecoin demand and parity deviations rise with global stress. The exchange rate OLS bias has structure analogous to the CIP case.

C Proofs

C.1 Proof of Proposition 1

Substituting market clearing $B^{sc} = \bar{\rho}^{sc} - \alpha_{sc} y$ and $B^{syn} = \bar{\rho}^{syn} + \tilde{\alpha}_\delta \delta$ into the first-order conditions (14)–(15):

$$-\delta = (\Gamma_\delta + \Gamma_\times)(\bar{\rho}^{syn} + \tilde{\alpha}_\delta \delta) + \Gamma_\times(\bar{\rho}^{sc} - \alpha_{sc} y)$$
$$y = \Gamma_\times(\bar{\rho}^{syn} + \tilde{\alpha}_\delta \delta) + (\Gamma_{sc} + \Gamma_\times)(\bar{\rho}^{sc} - \alpha_{sc} y)$$

Expanding and collecting terms in the first equation:

$$-\delta - \tilde{\alpha}_\delta(\Gamma_\delta + \Gamma_\times)\delta = (\Gamma_\delta + \Gamma_\times)\bar{\rho}^{syn} + \Gamma_\times \bar{\rho}^{sc} - \alpha_{sc}\Gamma_\times y$$
$$[1 + \tilde{\alpha}_\delta(\Gamma_\delta + \Gamma_\times)](-\delta) = (\Gamma_\delta + \Gamma_\times)\bar{\rho}^{syn} + \Gamma_\times \bar{\rho}^{sc} - \alpha_{sc}\Gamma_\times y$$

In the second equation:

$$y + \alpha_{sc}(\Gamma_{sc} + \Gamma_\times) y = (\Gamma_{sc} + \Gamma_\times)\bar{\rho}^{sc} + \Gamma_\times \bar{\rho}^{syn} + \tilde{\alpha}_\delta \Gamma_\times \delta$$

Using the definitions $A_\delta \equiv 1 + \tilde{\alpha}_\delta(\Gamma_\delta + \Gamma_\times)$, $A_{sc} \equiv 1 + \alpha_{sc}(\Gamma_{sc} + \Gamma_\times)$, $B \equiv \alpha_{sc}\Gamma_\times$, $C \equiv \tilde{\alpha}_\delta \Gamma_\times$, and writing $\delta = -(-\delta)$:

$$A_\delta(-\delta) + B\, y = (\Gamma_\delta + \Gamma_\times)\bar{\rho}^{syn} + \Gamma_\times \bar{\rho}^{sc} \equiv R_\delta$$
$$-C\,(-\delta) + A_{sc}\, y = (\Gamma_{sc} + \Gamma_\times)\bar{\rho}^{sc} + \Gamma_\times \bar{\rho}^{syn} \equiv R_{sc}$$

The determinant is $\Delta = A_\delta A_{sc} - BC > 0$. By Cramer's rule:

$$(-\delta)^* = \frac{A_{sc} R_\delta - B\, R_{sc}}{\Delta}, \qquad y^* = \frac{A_\delta R_{sc} + C\, R_\delta}{\Delta}$$

Expanding y^*, the coefficient of $\bar{\rho}^{sc}$ is $A_\delta(\Gamma_{sc} + \Gamma_\times) - \tilde{\alpha}_\delta \Gamma_\times^2 \equiv N_y$. The coefficient of $\bar{\rho}^{syn}$ simplifies to:

$$A_\delta \Gamma_\times - \tilde{\alpha}_\delta \Gamma_\times(\Gamma_\delta + \Gamma_\times) = \Gamma_\times[A_\delta - \tilde{\alpha}_\delta(\Gamma_\delta + \Gamma_\times)] = \Gamma_\times$$

Expanding $(-\delta)^*$, the coefficient of $\bar{\rho}^{syn}$ is $A_{sc}(\Gamma_\delta + \Gamma_\times) - \alpha_{sc}\Gamma_\times^2 \equiv N_\delta$. The coefficient of $\bar{\rho}^{sc}$ simplifies to:

$$A_{sc}\Gamma_\times - \alpha_{sc}\Gamma_\times(\Gamma_{sc} + \Gamma_\times) = \Gamma_\times[A_{sc} - \alpha_{sc}(\Gamma_{sc} + \Gamma_\times)] = \Gamma_\times$$

This yields equations (16)–(17). The exchange rate follows from (13). □

C.2 Proof of Proposition 2

Differentiating y^* with respect to $\bar{\rho}^{sc}$:

$$\frac{\partial y^*}{\partial \bar{\rho}^{sc}} = \frac{N_y}{\Delta} = \frac{A_\delta(\Gamma_{sc} + \Gamma_\times) - \tilde{\alpha}_\delta \Gamma_\times^2}{\Delta}$$

To verify $N_y > 0$, expand:

$$\begin{aligned} N_y &= [1 + \tilde{\alpha}_\delta(\Gamma_\delta + \Gamma_\times)](\Gamma_{sc} + \Gamma_\times) - \tilde{\alpha}_\delta \Gamma_\times^2 \\ &= (\Gamma_{sc} + \Gamma_\times) + \tilde{\alpha}_\delta \left[(\Gamma_\delta + \Gamma_\times)(\Gamma_{sc} + \Gamma_\times) - \Gamma_\times^2\right] \\ &= (\Gamma_{sc} + \Gamma_\times) + \tilde{\alpha}_\delta \left[\Gamma_\delta \Gamma_{sc} + \Gamma_\delta \Gamma_\times + \Gamma_\times \Gamma_{sc}\right] > 0 \end{aligned}$$

since all cost parameters and elasticities are non-negative. This proves Proposition 2 (i). □

Differentiating $(-\delta)^*$ with respect to $\bar{\rho}^{sc}$:

$$\frac{\partial (-\delta)^*}{\partial \bar{\rho}^{sc}} = \frac{\Gamma_\times}{\Delta}$$

so that:

$$\frac{\partial \delta^*}{\partial \bar{\rho}^{sc}} = -\frac{\Gamma_\times}{\Delta}$$

This is strictly negative when $\Gamma_\times > 0$ (since $\Delta > 0$) and zero when $\Gamma_\times = 0$, proving Proposition 2 (ii) and Corollary 1. □

For the exchange rate, differentiating (18) with respect to $\bar{\rho}^{sc}$:

$$\frac{\partial s^*}{\partial \bar{\rho}^{sc}} = \frac{1}{\nu_s}\left(-\eta_\delta \frac{\partial \delta^*}{\partial \bar{\rho}^{sc}} + \eta_y \frac{\partial y^*}{\partial \bar{\rho}^{sc}}\right) = \frac{1}{\nu_s}\left(\eta_\delta \frac{\Gamma_\times}{\Delta} + \eta_y \frac{N_y}{\Delta}\right) = \frac{\eta_\delta \Gamma_\times + \eta_y N_y}{\nu_s \Delta} > 0$$

since $\eta_\delta, \eta_y, \Gamma_\times, N_y, \nu_s, \Delta$ are all positive. This proves Proposition 2 (ii). □

Proof of Proposition 3

Since $\bar{\rho}_B^{sc} = \bar{D}_B^{sc} + S_B^{sc}$ and $\bar{D}_B^{sc}$ is unaffected by cross-book participant reallocation:

$$\frac{\partial \bar{\rho}_B^{sc}}{\partial S_B^{sc}} = 1$$

The results then follow directly from Proposition 2:

$$\frac{\partial y_B}{\partial S_B^{sc}} = \frac{\partial y^*}{\partial \bar{\rho}^{sc}} > 0, \qquad \frac{\partial \delta_B}{\partial S_B^{sc}} = \frac{\partial \delta^*}{\partial \bar{\rho}^{sc}} < 0, \qquad \frac{\partial s_B}{\partial S_B^{sc}} = \frac{\partial s^*}{\partial \bar{\rho}^{sc}} > 0$$

When S_B^{sc} falls due to reallocation toward country A, all effects reverse sign: country B's stablecoin parity deviation falls, its CIP deviation narrows (lower dollar premium), and its currency appreciates. □

C.3 Proof of Proposition 5

From the model with shock structure (38)–(43):

$$y_f = \frac{\partial y^{sc}}{\partial \bar{\rho}^{sc}}(\bar{D}_0 + \phi\mu + u_f - \tilde{S}_f^{sc}) + \xi_y\mu + \varepsilon_f^y$$
$$g_f = \gamma_g(\bar{D}_0 + \phi\mu + u_f - \tilde{S}_f^{sc}) - \alpha_{sc}(\xi_y\mu + \varepsilon_f^y)$$

The OLS coefficient is $\beta_y^{OLS} = \mathrm{Cov}(y_f, g_f)/\mathrm{Var}(g_f)$. Computing:

$$\mathrm{Cov}(y_f, g_f) = \frac{\partial y^{sc}}{\partial \bar{\rho}^{sc}}\gamma_g(\phi^2\sigma_\mu^2 + \sigma_u^2) + \phi\xi_y\gamma_g\sigma_\mu^2 - \alpha_{sc}\frac{\partial y^{sc}}{\partial \bar{\rho}^{sc}}(\xi_y\phi\sigma_\mu^2 + \sigma_{u\varepsilon}^y) + \text{other terms}$$

After simplification, the bias terms emerge as stated. □

C.4 Proof of Proposition 4

The GIV $Z_f = \sum_{f' \neq f} w_{f'}\tilde{u}_{f'}$ satisfies $\mathrm{Cov}(Z_f, \mu) = 0$ (excludes common shock), $\mathrm{Cov}(Z_f, u_f) = 0$ (excludes own shock), and $\mathrm{Cov}(Z_f, \varepsilon_f) = 0$ (Assumption 1(iii)).

Therefore:

$$\mathrm{Cov}(y_f, Z_f) = \frac{\partial y^{sc}}{\partial \bar{\rho}^{sc}} \cdot \mathrm{Cov}(\bar{\rho}_f^{sc}, Z_f)$$
$$\mathrm{Cov}(g_f, Z_f) = \gamma_g \cdot \mathrm{Cov}(\bar{\rho}_f^{sc}, Z_f)$$

where $\gamma_g \equiv 1 - \alpha_{sc}(\partial y^{sc}/\partial \bar{\rho}^{sc}) = 1 - \alpha_{sc}N_y/\Delta = (\Delta - \alpha_{sc}N_y)/\Delta$.

We show that $\Delta - \alpha_{sc}N_y = A_\delta$:

$$\begin{aligned}\Delta - \alpha_{sc}N_y &= A_\delta A_{sc} - BC - \alpha_{sc}\left[A_\delta(\Gamma_{sc} + \Gamma_\times) - \tilde{\alpha}_\delta\Gamma_\times^2\right] \\ &= A_\delta\left[A_{sc} - \alpha_{sc}(\Gamma_{sc} + \Gamma_\times)\right] - \alpha_{sc}\tilde{\alpha}_\delta\Gamma_\times^2 + \alpha_{sc}\tilde{\alpha}_\delta\Gamma_\times^2 \\ &= A_\delta \cdot 1 = A_\delta\end{aligned}$$

using $A_{sc} - \alpha_{sc}(\Gamma_{sc} + \Gamma_\times) = 1$ and $BC = \alpha_{sc}\tilde{\alpha}_\delta\Gamma_\times^2$.

Therefore $\gamma_g = A_\delta/\Delta$, and the IV coefficients are:

Part (i).

$$\beta_y^{IV} = \frac{\partial y^{sc}/\partial \bar{\rho}^{sc}}{\gamma_g} = \frac{N_y/\Delta}{A_\delta/\Delta} = \frac{N_y}{A_\delta}$$

Part (ii).

$$\beta_\delta^{IV} = \frac{\partial \delta^{sc}/\partial \bar{\rho}^{sc}}{\gamma_g} = \frac{-\Gamma_\times/\Delta}{A_\delta/\Delta} = -\frac{\Gamma_\times}{A_\delta}$$

Part (iii). Follows from the exchange rate equation and parts (i)–(ii).

Part (iv). The ratio:

$$\frac{|\beta_\delta^{IV}|}{\beta_y^{IV}} = \frac{\Gamma_\times/A_\delta}{N_y/A_\delta} = \frac{\Gamma_\times}{N_y}$$

To verify $\Gamma_\times/N_y < 1$, note that:

$$N_y = (\Gamma_{sc} + \Gamma_\times) + \tilde{\alpha}_\delta(\Gamma_\delta\Gamma_{sc} + \Gamma_\delta\Gamma_\times + \Gamma_\times\Gamma_{sc}) > \Gamma_\times$$

since $\Gamma_{sc} > 0$, so $\Gamma_\times/N_y < 1$. The ratio is strictly positive when $\Gamma_\times > 0$ and zero when $\Gamma_\times = 0$. □

D Alternative model specification: Segmented households

This appendix develops an alternative specification in which households are segmented into two types: traditional households who access only the spot and swap market, and crypto-native households who access only the stablecoin market. This specification may be appropriate when institutional barriers prevent certain households from accessing certain markets.

D.1 Setup

Traditional households. Traditional households (mass μ_T) have access to spot FX and FX swaps but cannot hold stablecoins. This may reflect regulatory restrictions on cryptocurrency holdings in certain jurisdictions, lack of access to crypto infrastructure (wallets, exchanges, on-ramps), institutional mandates that prohibit digital asset exposure, or simple unfamiliarity with blockchain technology. These households demand synthetic dollars for hedging, trade finance and treasury management:

$$D^T = \bar{D}^T + \alpha_T \delta \tag{48}$$

where $\bar{D}^T$ captures baseline hedging needs and $\alpha_T > 0$ is the demand elasticity with respect to the CIP deviation. When δ falls (synthetic dollars become more expensive), demand for synthetic dollars decreases.

Crypto-native households. Crypto-native households (mass μ_C) have access to spot FX and stablecoins but cannot access swap markets. This may reflect lack of banking relationships required for swap market participation, failure to meet minimum size requirements for institutional FX markets, exclusion from traditional financial infrastructure due to informal employment or lack of credit history, or preference for decentralized financial infrastructure. These households demand stablecoins for remittances, e-commerce, capital preservation, and participation in decentralized finance:

$$D^C = \bar{D}^C - \alpha_C y \tag{49}$$

where $\bar{D}^C$ captures baseline demand for on-chain dollar access and α_C is the demand elasticity with respect to the stablecoin parity deviation.

D.2 Market clearing

The stablecoin market clears when crypto-native demand equals relative supply plus intermediary provision:

$$D^C = \tilde{S}^{sc} + B^{sc} \tag{50}$$

Substituting the demand function and defining $\bar{\rho}^{sc} \equiv \bar{D}^C - \tilde{S}^{sc}$:

$$B^{sc} = \bar{\rho}^{sc} - \alpha_C y \tag{51}$$

The swap market clears when traditional demand equals US supply plus intermediary provision:

$$D^T = L^{US} + B^{syn} \tag{52}$$

Defining $\bar{\rho}^{syn} \equiv \bar{D}^T - \bar{L}$ and $\tilde{\alpha}_\delta \equiv \alpha_T + \alpha_L$:

$$B^{syn} = \bar{\rho}^{syn} + \tilde{\alpha}_\delta \delta \tag{53}$$

The spot market clears as before, with both household types participating:

$$Q^{spot} = \bar{Q} - \eta_\delta \delta + \eta_y y - \nu_s s = 0 \tag{54}$$

D.3 Equilibrium

With the same intermediary cost function as in the main text, the intermediary's first-order conditions are unchanged. The equilibrium therefore has identical mathematical structure:

$$\delta^* = \frac{A_{sc} R_\delta - B R_{sc}}{\Delta} \tag{55}$$

$$y^* = \frac{A_\delta R_{sc} + C R_\delta}{\Delta} \tag{56}$$

where now $\alpha_{sc} = \alpha_C$, $\tilde{\alpha}_\delta = \alpha_T + \alpha_L$, and all composite parameters are defined as before with these substitutions.

All comparative statics results from Section 4 carry over. In particular, stablecoin demand shocks (now interpreted as shocks to crypto-native demand $\bar{D}^C$) raise stablecoin parity deviations, lower CIP deviations (increasing the dollar premium) when $\Gamma_\times > 0$, and depreciate the local currency.

D.4 Equivalence of reduced forms

The segmented and integrated specification in the main text yield identical reduced-form equilibrium conditions. The demand functions have the same linear form, and the intermediary's problem is unchanged. Consequently, the equilibrium prices (δ^*, y^*, s^*) respond identically to demand shocks in both specifications.

In the integrated specification, a shock to $\bar{D}^{sc}$ represents a change in the representative household's demand for stablecoins (because of e.g. increased use of stablecoins for remittances or a shift toward crypto-based commerce). In the segmented specification, a shock to $\bar{D}^C$ represents either a change in crypto-native households' demand intensity or a change in the mass μ_C of crypto-native households (e.g., due to broader adoption of cryptocurrency).

In the integrated specification, households substitute between stablecoins and synthetic dollars in response to relative price changes. If stablecoin parity deviations rise, households shift toward swaps; if CIP falls (dollar premium rises), households shift toward stablecoins. This substitution dampens price movements and provides a form of market integration even absent intermediary linkages.

In the segmented specification, no such substitution occurs. Traditional households cannot shift to stablecoins when CIP falls; crypto-native households cannot shift to swaps when stablecoin parity deviations rise. Price movements are therefore larger for a given shock, and the markets are more fragile.

The choice between specifications should depend on the empirical context. The integrated specification presented in the main text may be more appropriate for developed markets or for institutional investors who can access multiple channels. Large corporations, for instance, can

choose between hedging via FX swaps and holding stablecoins for treasury management. For these agents, the two channels are genuine substitutes.

The segmented specification may be more appropriate for emerging markets where institutional barriers are more binding, or for retail investors who face significant frictions in accessing traditional FX markets. In these contexts, the assumption that households are confined to specific market segments may be more realistic.

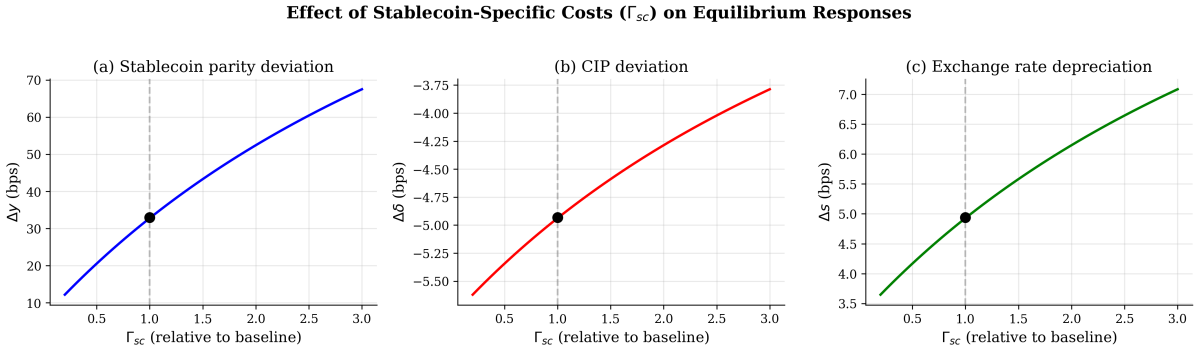
E Additional counterfactuals

Stablecoin-specific costs. Figure 9 reveals that stablecoin-specific costs primarily affect the own-market response rather than the spillover. Specifically, doubling Γ_{sc} more than doubles the elasticity of parity deviations. However, the effect on CIP is muted.

This asymmetry arises because Γ_{sc} affects only stablecoin intermediation and only affecting CIP through the equilibrium balance sheet effects. Higher values of Γ_{sc} imply that intermediaries require a larger premium to absorb a given level of stablecoin demand, which amplifies Δy. The modest decline in the dollar premium reflects that larger stablecoin parity deviations induce some substitution toward swaps. This suggests that policies targeting stablecoin market efficiency may not substantially reduce systemic risk transmission to traditional markets.

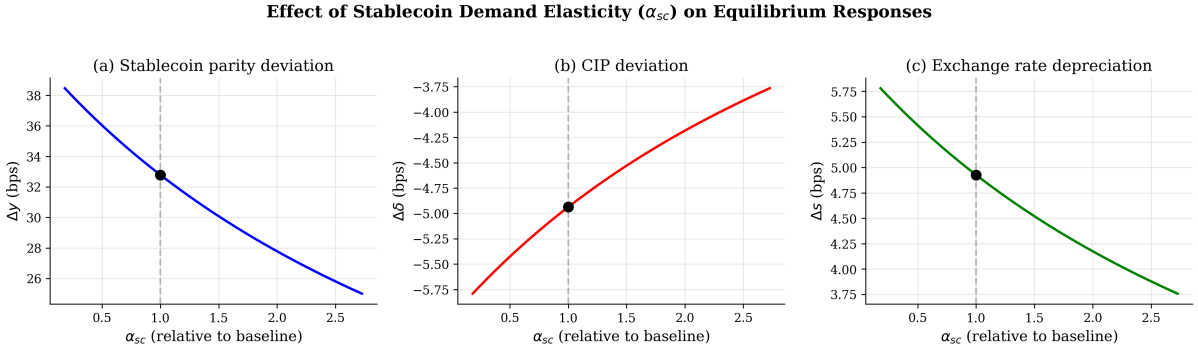
Stablecoin demand elasticity. Figure 10 shows the equilibrium responses under alternative demand elasticities. Counterfactuals show larger responses for lower elasticities, which approach to the identified supply responses in the limit to fully inelastic demand. Nevertheless, the range of parity deviations response goes from 40 to 26 under different parameterizations. Spillovers are also in a tight range.

Figure 9: Effect of stablecoin-specific costs (Γ_{sc}) on price responses.



Notes: This figure displays equilibrium price responses to a 1% stablecoin flow shock across different stablecoin-specific cost parameter Γ_{sc}. The left panel shows the stablecoin parity deviation (Δy), the middle panel shows the CIP deviation ($\Delta\delta$, where more negative values indicate a higher dollar premium), and the right panel shows the exchange rate depreciation (Δs). The horizontal axis measures Γ_{sc} relative to its calibrated baseline value. Dots indicate baseline values.

Figure 10: Effect of stablecoin demand elasticity (α_{sc}) on Price Responses.



Notes: This figure displays equilibrium price responses to a 1% stablecoin flow shock across different stablecoin demand elasticity α_{sc}. The left panel shows the stablecoin parity deviation (Δy), the middle panel shows the CIP deviation ($\Delta\delta$, where more negative values indicate a higher dollar premium), and the right panel shows the exchange rate depreciation (Δs). The horizontal axis measures α_{sc} relative to its calibrated baseline value. Dots indicate baseline values.

F Dynamic model extension

This appendix extends the static model to incorporate persistence in demand shocks and intermediary balance sheet dynamics. The extension draws on the intermediary asset pricing literature, particularly He and Krishnamurthy (2013), who demonstrate that when financial intermediaries face equity capital constraints, risk premia become state-dependent and exhibit nonlinear dynamics during crises. In the international finance context, Gabaix and Maggiori (2015) show that capital flows drive exchange rates by altering the balance sheets of financiers who bear currency risk, while Du et al. (2018) document that CIP deviations reflect the shadow cost of intermediary balance sheet capacity. Our dynamic extension captures these mechanisms in a tractable framework that generates persistent spillovers and state-dependent amplification. We develop the full extension here, and refer to Figures reproduced in the main text.

F.1 Environment

Time is discrete and indexed by $t = 0, 1, 2, \ldots$. The structure of markets and agents follows the static model, with two key additions: demand shocks exhibit persistence, and intermediary risk-bearing capacity varies with wealth.

Demand dynamics. Stablecoin demand follows an AR(1) process:

$$\bar{D}_t^{sc} = \rho \bar{D}_{t-1}^{sc} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2) \tag{57}$$

where $\rho \in (0, 1)$ governs persistence. This specification captures the empirical observation that stablecoin flows exhibit positive autocorrelation, with demand shocks decaying gradually rather than reverting immediately. We calibrate $\rho = 0.8$, implying a half-life of approximately three days, consistent with the observed persistence in stablecoin trading volumes.

Intermediary problem with wealth-dependent risk capacity. The representative intermediary enters period t with wealth W_t and chooses positions in synthetic dollar intermediation (B_t^{syn}) and stablecoin intermediation (B_t^{sc}). The key departure from the static model is that the intermediary's costs are scaled by wealth, capturing the idea that positions must be evaluated relative to capital. The intermediary solves:

$$\max_{B_t^{syn}, B_t^{sc}} \quad \delta_t B_t^{syn} + y_t B_t^{sc} - \frac{\Gamma_\delta}{2W_t}(B_t^{syn})^2 - \frac{\Gamma_{sc}}{2W_t}(B_t^{sc})^2 - \frac{\Gamma_\times}{2W_t}(B_t^{syn} + B_t^{sc})^2 \tag{58}$$

The scaling by W_t in the cost terms admits a natural interpretation: a given dollar position represents greater risk exposure when measured against a smaller capital base. An intermediary

with \$10 billion in capital can comfortably absorb losses on a \$1 billion position; the same position would be precarious for an intermediary with only \$2 billion. This formulation is analogous to measuring portfolio risk relative to equity, as in the VaR-constraint literature (Brunnermeier and Pedersen, 2009; Adrian and Shin, 2014), and captures the risk-bearing capacity mechanism emphasized by Gabaix and Maggiori (2015).

The first-order conditions yield optimal positions that scale linearly with wealth:

$$B_t^{syn} = W_t \cdot b^{syn}(\delta_t, y_t; \Gamma) \tag{59}$$

$$B_t^{sc} = W_t \cdot b^{sc}(\delta_t, y_t; \Gamma) \tag{60}$$

where $b^{syn}(\cdot)$ and $b^{sc}(\cdot)$ are the position-per-unit-wealth functions derived from the static model's first-order conditions. When wealth falls, the intermediary endogenously reduces positions, withdrawing risk-bearing capacity from both markets.

Wealth dynamics. Intermediary wealth evolves with portfolio returns:

$$W_{t+1} = W_t \left(1 + \frac{B_t^{syn}}{W_t}\delta_t + \frac{B_t^{sc}}{W_t} y_t - \phi \cdot \Delta s_t\right) = W_t \left(1 + b_t^{syn}\delta_t + b_t^{sc} y_t - \phi \cdot \Delta s_t\right) \tag{61}$$

where ϕ captures the intermediary's net foreign exchange exposure and Δs_t is the rate of exchange rate depreciation. The position-to-wealth ratios b_t^{syn} and b_t^{sc} determine how portfolio returns translate into wealth growth.

This specification creates scope for amplification. Adverse shocks that raise the stablecoin parity deviation y_t and lower δ_t (increasing the dollar premium) may generate mark-to-market losses that deplete wealth. With reduced wealth, the intermediary's risk-bearing capacity falls, positions contract, and prices must adjust further to clear markets. This feedback loop is the central mechanism in He and Krishnamurthy (2013): losses reduce capacity, and reduced capacity sustains elevated prices.

F.2 Equilibrium

Within each period, given the state variables $(W_t, \bar{D}_t^{sc})$, equilibrium prices clear markets with intermediary positions scaling with wealth. Market clearing requires:

$$W_t \cdot b^{syn}(\delta_t, y_t) = \bar{\rho}^{syn} + \tilde{\alpha}_\delta \delta_t \tag{62}$$

$$W_t \cdot b^{sc}(\delta_t, y_t) = \bar{\rho}_t^{sc} - \alpha_{sc} y_t \tag{63}$$

where $\bar{\rho}_t^{sc}$ incorporates the time-varying demand shock. These conditions implicitly define equilibrium prices as functions of the state:

$$y_t = f_y(W_t, \bar{D}_t^{sc}; \Gamma) \tag{64}$$

$$\delta_t = f_\delta(W_t, \bar{D}_t^{sc}; \Gamma) \tag{65}$$

$$s_t = f_s(W_t, \bar{D}_t^{sc}; \Gamma) \tag{66}$$

The key comparative static are $\partial y_t / \partial W_t < 0$ and $\partial \delta_t / \partial W_t > 0$: lower intermediary wealth implies higher stablecoin parity deviations and lower CIP deviations (a higher dollar premium). Intuitively, when the intermediary has less capital, it supplies less risk-bearing capacity to the market. With reduced supply of intermediation, prices must adjust to equilibrate demand.

To characterize the equilibrium more explicitly, we normalize steady-state wealth to $\bar{W} = 1$. At steady state, the model reduces to the static framework analyzed in the main text. Away from steady state, we can write the equilibrium price functions as:

$$y_t = \frac{1}{W_t} \cdot \tilde{f}_y(\bar{D}_t^{sc}; \Gamma) \tag{67}$$

$$\delta_t = W_t \cdot \tilde{f}_\delta(\bar{D}_t^{sc}; \Gamma) \tag{68}$$

where $\tilde{f}_y$ and $\tilde{f}_\delta$ are the static equilibrium functions. The stablecoin parity deviation scales inversely with wealth, while the CIP deviation scales positively: a 50 percent reduction in intermediary capital doubles the stablecoin parity deviation and halves the CIP deviation (doubling the dollar premium in absolute terms) for a given demand shock.

Proposition 6 (Dynamic amplification)**.** *Consider an impulse response to a demand shock $\varepsilon_0 > 0$ from steady state ($W_0 = \bar{W}, \bar{D}_0^{sc} = 0$). Let $\mathcal{A}_y \equiv \sum_{t=0}^{\infty} y_t / y_0$ denote the cumulative amplification factor for the stablecoin parity deviation. Then:*

(i) $\mathcal{A}_y > 1/(1-\rho)$ when the balance sheet channel is active, exceeding the amplification from demand persistence alone.

(ii) $\partial \mathcal{A}_y / \partial W_0 < 0$: lower initial wealth increases cumulative amplification.

(iii) Analogous results hold for the CIP deviation δ and exchange rate s.

The proof follows from iterating the equilibrium conditions forward. With $\rho > 0$, demand shocks persist, sustaining elevated prices. Simultaneously, if the shock depletes wealth ($W_1 < W_0$), reduced risk-bearing capacity amplifies the price response in subsequent periods. The two channels compound: persistent demand meets diminished intermediation capacity, generating cumulative effects that exceed what either channel would produce in isolation.

F.3 Calibration

We calibrate the dynamic model to match the static calibration at steady state, adding parameters governing dynamics. Table 6 reports the parameter values.

Parameter	Description	Value	Source
Static parameters (from main calibration)			
Γ_δ	Swap market cost	0.023	CIP literature
Γ_{sc}	Stablecoin market cost	0.194	Match Δy
$\Gamma_\times$	Cross-market cost	0.148	Match $\Delta\delta$
Dynamic parameters			
ρ	Demand persistence	0.80	Stablecoin flow autocorrelation
ϕ	FX exposure	150	Intermediary hedging ratios
$\bar{W}$	Steady-state wealth	1	Normalization

Table 6: Dynamic model calibration. Static parameters are taken from the main text. With steady-state wealth normalized to unity, the static model emerges as the $W_t = 1$ case. Dynamic parameters are chosen to match empirical features of stablecoin markets and intermediary balance sheets.

The persistence parameter $\rho = 0.8$ implies that demand shocks have a half-life of approximately three days. The FX exposure parameter ϕ governs how exchange rate movements feed back into intermediary wealth; we calibrate it to match evidence on intermediary hedging practices in foreign exchange markets.

F.4 Impulse response analysis

Figure 7 displays the impulse response to a one percent stablecoin demand shock, starting from steady state. The shock generates impact effects of 34 basis points for the stablecoin parity deviation y, -15 basis points for the CIP deviation δ (a higher dollar premium), and 3.5 basis points for exchange rate depreciation s. These impact effects coincide with the static model, as the economy begins at steady-state wealth.

The dynamic responses reveal the interaction between demand persistence and balance sheet effects. On impact, intermediary wealth drops by approximately 5 percent as mark-to-market losses materialize. This wealth depletion reduces risk-bearing capacity in subsequent periods, amplifying the price response relative to what demand persistence alone would generate. Prices decay gradually, with a half-life of roughly four days, somewhat longer than the three-day half-life implied by the AR(1) demand process. The additional persistence reflects the balance sheet channel: even as demand shocks fade, depleted wealth sustains elevated stablecoin parity deviations and depressed CIP deviations.

Cumulating the impulse responses yields amplification factors of 5.4 for the stablecoin parity deviation and 5.6 for the CIP deviation. These figures indicate that static analysis, which

captures only impact effects, understates the true spillover costs by a factor of five to six. Demand persistence alone would generate amplification of $1/(1-\rho) = 5$; the additional 8 to 12 percent comes from the wealth channel.

F.5 State-dependent amplification

A key prediction of the model is that identical shocks generate larger price responses when intermediary balance sheets are impaired. This property follows directly from the equilibrium structure: with the stablecoin parity deviation scaling as $1/W_t$ and the CIP deviation scaling as W_t, lower wealth translates mechanically into higher stablecoin parity deviations and larger dollar premia. Table 7 reports impact effects for the same demand shock under alternative assumptions about initial wealth. The table reports price responses to a one percent demand shock under alternative assumptions about initial intermediary wealth. Negative values of $\Delta\delta$ indicate a higher dollar premium. Amplification is computed relative to the steady-state baseline.

Initial Wealth	Δy (bps)	$\Delta\delta$ (bps)	Δs (bps)	Amplification
$W_0 = \bar{W}$ (steady state)	34	−15	3.5	1.0×
$W_0 = 0.75\bar{W}$	45	−20	4.7	1.3×
$W_0 = 0.50\bar{W}$ (stressed)	68	−30	7.0	2.0×

Table 7: State-dependent impact effects

When intermediary wealth is 50 percent below steady state, impact effects double relative to the baseline. This proportionality (stablecoin parity deviations scaling inversely with wealth, dollar premia scaling inversely with wealth) is a direct consequence of the risk-bearing capacity formulation. The economic content is that a given demand shock requires more price adjustment to clear markets when intermediation capacity is scarce.

The state dependence has important implications for crisis dynamics. An initial shock depletes intermediary capital, reducing risk-bearing capacity. Subsequent shocks then encounter a more constrained intermediary sector and generate amplified price responses, which may cause further wealth depletion. This feedback mechanism can produce destabilizing dynamics in which crises become self-reinforcing, as in He and Krishnamurthy (2013).

F.6 Counterfactual: Cross-market frictions

The static analysis identified the cross-market cost $\Gamma_\times$ as the key parameter governing spillovers from stablecoin markets to CIP deviations. The dynamic model amplifies this role through the balance sheet channel. Figure 8 compares impulse responses under the baseline calibration to a counterfactual with doubled cross-market costs.

Doubling $\Gamma_\times$ increases the impact effect on the stablecoin parity deviation from 34 to 78 basis points, a factor of 2.3. The CIP spillover increases from -15 to -58 basis points, nearly a fourfold increase in the dollar premium. The more-than-proportional increase in CIP spillovers reflects compounding through the balance sheet channel: higher $\Gamma_\times$ generates larger initial price responses, which cause greater wealth depletion, which further amplifies subsequent price dynamics.

These findings reinforce the policy implications from the static analysis. Cross-market frictions, whether arising from balance sheet linkages, regulatory constraints, or market segmentation, are the primary determinant of spillover magnitude. The dynamic model reveals that reducing these frictions yields compounding benefits: smaller spillovers preserve intermediary capital, which sustains risk-bearing capacity, which further dampens price responses to future shocks.

F.7 Discussion

The dynamic extension connects our analysis of stablecoin-FX spillovers to several strands of the literature on financial frictions and international finance.

The model provides a specific application of the intermediary asset pricing framework developed by He and Krishnamurthy (2013). In their setting, intermediaries face an equity capital constraint that limits positions to a multiple of wealth. When binding, this constraint creates state-dependent risk premia: intermediation capacity becomes scarce, and the marginal intermediary demands higher compensation for bearing risk. Our formulation captures the same economic force through a different modeling device (wealth-scaled costs rather than a hard constraint) but delivers the same qualitative predictions. Risk premia rise when intermediary wealth falls, and the amplification is nonlinear, with crises generating disproportionately large effects.

The framework also speaks to the literature on CIP deviations and limits to arbitrage. Du et al. (2018) document that CIP violations reflect binding balance sheet constraints, particularly around regulatory reporting dates when banks face pressure to reduce positions. In their interpretation, the cross-currency basis represents the shadow value of balance sheet space: arbitrageurs could eliminate CIP deviations, but doing so requires capital that is costly to deploy. Our model formalizes this mechanism. The term Γ/W_t in the intermediary's cost function captures exactly this shadow value, the effective price of deploying a unit of balance sheet capacity. When wealth is depleted, this shadow price rises, and CIP deviations persist even in the presence of apparent arbitrage opportunities.

Our analysis complements Gabaix and Maggiori (2015), who model exchange rate determination through the balance sheets of financiers intermediating international capital flows. In their framework, capital flows alter intermediary wealth and thereby affect the compensation required

for bearing currency risk. Our contribution is to incorporate stablecoins as an additional market and trace how shocks propagate to traditional FX pricing through shared intermediary balance sheets. The risk-bearing capacity formulation we adopt is closely related to their approach, with the Γ parameter playing an analogous role in governing how balance sheet conditions translate into price dynamics.

Finally, the state-dependent amplification documented here has implications for financial stability. The potential for destabilizing feedback, in which shocks deplete capital, reduce intermediation capacity, amplify price responses, and further deplete capital, creates fragility in the system. As stablecoin markets grow and their integration with traditinal finance deepens, this feedback channel may become increasingly important. Prudential policies targeting intermediary capital buffers could dampen the amplification mechanism, though at the cost of reduced intermediation in normal times.

F.8 Redemption frictions and run dynamics

We extend the model to incorporate redemption frictions that amplify during stress, providing a framework for analyzing run dynamics in stablecoin markets. The key insight is that large redemptions relative to intermediary capacity impose additional costs, creating a feedback loop that can destabilize the system.

Haircut structure. We introduce a *redemption haircut* $h_t \in (0, 1]$ that reduces the effective value recovered by the intermediary on each unit of stablecoin position during large redemption episodes. Specifically, h_t captures the combination of fire-sale discounts (reserve assets sold at below-par prices to meet redemptions), settlement frictions (delays in converting illiquid reserves to cash), and issuer liquidity risk (the risk that the stablecoin issuer cannot honor redemptions at par). The haircut is borne by the intermediary who holds the stablecoin position B_t^{sc}: when households redeem, the intermediary receives only h_t dollars per dollar of face value. We parameterize it as:

$$h_t = 1 - \eta \frac{|B_t^{sc}|}{W_t} \tag{69}$$

where $\eta \geq 0$ governs the severity of redemption frictions and $|B_t^{sc}|/W_t$ is the redemption volume scaled by intermediary capital. When $\eta = 0$, we recover the baseline model with no redemption costs. When $\eta > 0$, larger redemption volumes relative to the intermediary's capital base imply steeper haircuts, reflecting that forced asset sales become costlier as they grow relative to the intermediary's ability to absorb losses.

Modified wealth dynamics. Intermediary wealth evolves as:

$$W_{t+1} = W_t \left(1 + b_t^{syn} \delta_t + b_t^{sc} y_t \cdot h_t - \phi \Delta s_t\right) \tag{70}$$

The haircut reduces the effective return on stablecoin intermediation. During redemption episodes ($b_t^{sc} < 0$), the intermediary faces losses that are amplified by the factor $(2 - h_t)$: as the haircut rises (lower h_t), losses mount nonlinearly.

Run threshold. The solvency constraint $h_t > 0$ defines a critical threshold. Substituting (69) and using $|b_t^{sc}| = |B_t^{sc}|/W_t$, the constraint becomes:

$$\eta \cdot |b_t^{sc}| < 1 \quad \Longleftrightarrow \quad \eta < \frac{W_t}{|B_t^{sc}|} \tag{71}$$

For a given shock that generates position $|B_t^{sc}|$, the *run threshold* is:

$$\bar{\eta} = \frac{W_t}{\text{shock size}} \tag{72}$$

When redemption frictions exceed this threshold, the haircut turns negative, the intermediary cannot honor redemptions at any positive value, and the system breaks down.

Feedback mechanism. The interaction between haircuts and wealth creates a destabilizing feedback loop:

1. A redemption shock increases $|b_t^{sc}|$
2. The haircut rises (lower h_t), reducing value recovered
3. Wealth W_{t+1} falls more than in the baseline model
4. Lower wealth raises effective costs (Γ/W_{t+1}), widening spreads
5. The run threshold $\bar{\eta} = W_{t+1}/\text{shock}$ falls, making the system more fragile

This mechanism captures the classic ingredients of a bank run applied to stablecoin intermediation: sequential service constraints (early redeemers get better terms), strategic complementarities (my incentive to redeem rises if others redeem), and the possibility of self-fulfilling crises.

Counterfactual analysis. Figure 11 illustrates how redemption frictions amplify price responses for different levels of intermediary capitalization. Each curve shows the impact effect of a 1% redemption shock as a function of η, with vertical dotted lines marking the run threshold $\bar{\eta} = W_0$ for each case.

Three features emerge from the figure. First, *amplification is convex*: price responses increase slowly for low η but accelerate sharply as frictions approach the threshold. At $\eta = 0.9\bar{\eta}$, impact effects are roughly ten times larger than in the frictionless baseline. Second, *capitalization determines fragility*: a well-capitalized intermediary ($W_0 = 1$) can tolerate frictions up to $\eta = 1$, while a stressed intermediary ($W_0 = 0.5$) hits the run region at $\eta = 0.5$. The same level of reserve

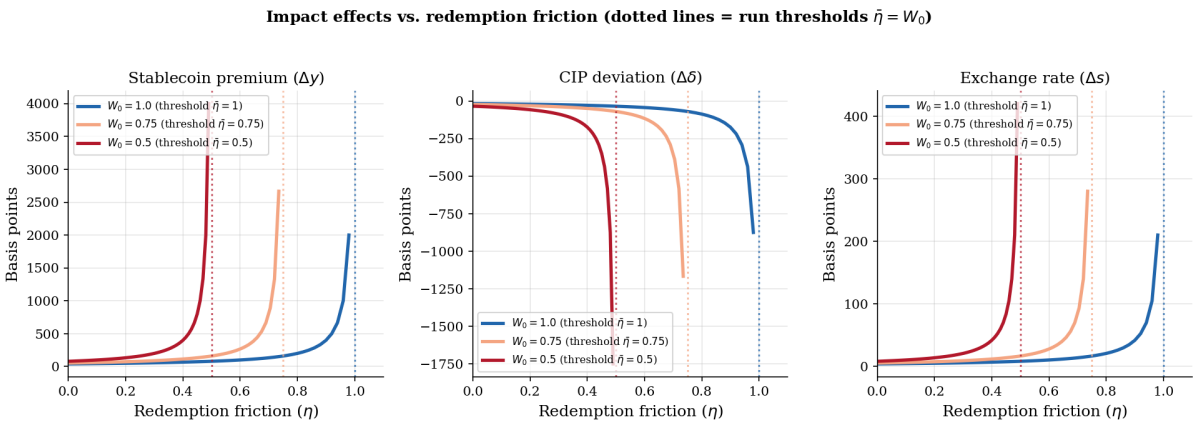


Figure 11: Effect of redemption frictions (η) on impact responses to a 1% stablecoin redemption shock. Panels show the stablecoin parity deviation (left), CIP deviation (middle), and exchange rate depreciation (right). Each curve corresponds to a different initial wealth level W_0. Dotted vertical lines mark the run threshold $\bar{\eta} = W_0$ for each case; effects diverge to infinity as η approaches the threshold.

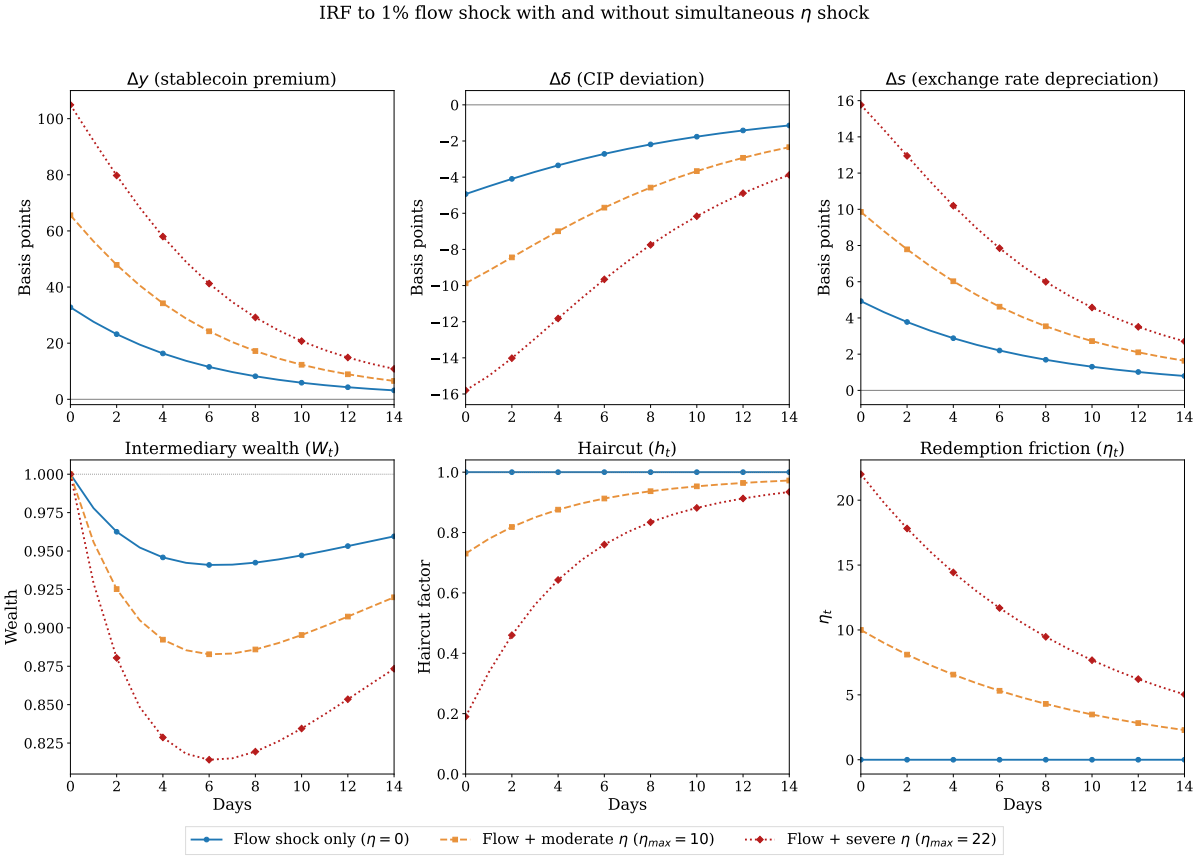
illiquidity that is manageable in normal times becomes catastrophic when balance sheets are impaired. Third, *spillovers scale proportionally*: the ratio of CIP response to stablecoin parity deviation response remains approximately constant as η varies, implying that the cross-market spillover channel operates similarly across the friction spectrum.

Policy implications. The analysis yields two policy-relevant insights. First, *reserve liquidity requirements* (reducing η) provide a buffer against runs by keeping the system away from the threshold. Mandating that stablecoin issuers hold liquid reserves,such as Treasury bills rather than commercial paper or crypto assets,directly lowers η and expands the stable region. Second, *intermediary capital requirements* (raising W_t) increase the run threshold $\bar{\eta}$, allowing the system to tolerate higher frictions without destabilizing. The two policies are complementary: liquid reserves reduce the severity of redemption costs conditional on a run, while adequate capital reduces the probability of approaching run conditions in the first place.

Combined shocks and crisis amplification. The amplifying role of redemption frictions is most apparent when flow shocks and friction shocks occur simultaneously—as would be expected during a crisis episode. Figure 12 illustrates this by comparing the impulse response to a 1% flow shock under three scenarios: no friction ($\eta = 0$), moderate friction ($\eta = 0.4$), and severe friction ($\eta = 0.8$). In the frictionless baseline, the flow shock generates a 40 basis point increase in the stablecoin parity deviation and a 17.5 basis point widening of the CIP deviation, consistent with our empirical estimates. When the same flow shock is accompanied by a moderate friction shock, impact effects rise by a factor of 1.7; under severe frictions, the amplification reaches 5 times the baseline. This pattern captures the dynamics of a crisis episode such as the Terra/Luna collapse or the SVB-induced USDC depeg: bad

news triggers redemptions while simultaneously raising the cost of those redemptions, as counterparties become cautious, liquidity evaporates, and reserve assets face fire-sale discounts. The interaction between flow and friction shocks is central to understanding why stablecoin stress events can generate outsized spillovers to traditional FX markets.

Figure 12: Impulse response to a 1% stablecoin flow shock with and without a simultaneous shock to redemption frictions η.



Notes: The blue solid line shows the baseline response ($\eta = 0$); the orange dashed line adds a moderate friction shock; the red dotted line adds a severe friction shock. Upper panels show price responses (stablecoin parity deviation, CIP deviation, exchange rate); lower panels show intermediary wealth, the haircut factor, and the friction path. Both shocks decay over time.

G Finite supply elasticity

This appendix shows that relaxing the perfectly elastic supply assumption does not affect the identified structural object. We consider a model without cross-book participants but with finite supply elasticity and household demand that depends on the absolute stablecoin price.

G.1 Setup

Stablecoin supply responds to the dollar price deviation $\mu \equiv P_t^{s,\mathrm{USD}} - 1$:

$$S^{sc} = S_0 + \kappa\,\mu \tag{73}$$

Household demand in country f depends on the absolute price $P_f = 1 + \mu + y_f$:

$$D_f^{sc} = \bar{D}_f^{sc} - \alpha_{sc}(1 + \mu + y_f) \tag{74}$$

All other elements — the intermediary cost function, swap market, and spot market — are as in the main text.

G.2 Local equilibrium given μ

Define effective local demand $\bar{\rho}_f^{sc}(\mu) \equiv \tilde{D}_f^{sc} - \alpha_{sc}\mu$, where $\tilde{D}_f^{sc} \equiv \bar{D}_f^{sc} - \alpha_{sc}$. Market clearing on book f requires $B_f^{sc} = \bar{\rho}_f^{sc}(\mu) - \alpha_{sc}\, y_f$. Substituting into the intermediary's first-order conditions and solving yields:

$$y_f = \frac{N_y\,\bar{\rho}_f^{sc}(\mu) + N_{syn}\,\bar{\rho}^{syn}}{\Delta} \tag{75}$$

$$-\delta_f = \frac{\Gamma_\times\,\bar{\rho}_f^{sc}(\mu) + M_\delta\,\bar{\rho}^{syn}}{\Delta} \tag{76}$$

where N_y, N_{syn}, M_δ, and Δ are as defined in the main text.

G.3 Global clearing

The aggregate condition $\sum_f B_f^{sc} + D_{US} = S_0 + \kappa\mu$ pins μ as a function of all demand shocks $\{\tilde{D}_f^{sc}\}$.

G.4 Cross-country transmission and IV coefficient

An idiosyncratic shock u_A in country A ($\tilde{D}_A^{sc} \to \tilde{D}_A^{sc} + u_A$) affects country B by raising μ. With $\tilde{D}_B^{sc}$ unchanged, let $m \equiv \partial\mu/\partial u_A > 0$:

$$\frac{\partial y_B}{\partial u_A} = -\frac{N_y\,\alpha_{sc}}{\Delta} \cdot m \tag{77}$$

$$\frac{\partial B_B^{sc}}{\partial u_A} = -\alpha_{sc}\,m - \alpha_{sc}\frac{\partial y_B}{\partial u_A} \tag{78}$$

Taking the ratio:

$$\beta_y^{IV} = \frac{\partial y_B / \partial u_A}{\partial B_B^{sc} / \partial u_A} = \frac{-N_y \alpha_{sc}\, m/\Delta}{-\alpha_{sc}\, m(1 + N_y \alpha_{sc}/\Delta)} \tag{79}$$

The factor m cancels from numerator and denominator:

$$\beta_y^{IV} = \frac{N_y/\Delta}{(\Delta + N_y \alpha_{sc})/\Delta} = \frac{N_y}{\Delta - \alpha_{sc} N_y} \tag{80}$$

Using the identity $\Delta - \alpha_{sc} N_y = A_\delta$ (proved below):

$$\beta_y^{IV} = \frac{N_y}{A_\delta} \tag{81}$$

This is **identical** to the coefficient identified under the cross-book participant specification in the main text. The result holds for any $\kappa \in (0, \infty)$: the supply elasticity does not enter the identified object.

Proof that $\Delta - \alpha_{sc} N_y = A_\delta$. Recall:

$$\begin{aligned}
\Delta &= A_\delta A_{sc} - BC \\
N_y &= A_\delta(\Gamma_{sc} + \Gamma_\times) - \tilde{\alpha}_\delta \Gamma_\times^2 \\
A_{sc} &= 1 + \alpha_{sc}(\Gamma_{sc} + \Gamma_\times), \quad B = \alpha_{sc}\Gamma_\times, \quad C = \tilde{\alpha}_\delta \Gamma_\times
\end{aligned}$$

Then:

$$\begin{aligned}
\Delta - \alpha_{sc} N_y &= A_\delta A_{sc} - BC - \alpha_{sc}[A_\delta(\Gamma_{sc} + \Gamma_\times) - \tilde{\alpha}_\delta \Gamma_\times^2] \\
&= A_\delta[A_{sc} - \alpha_{sc}(\Gamma_{sc} + \Gamma_\times)] - BC + \alpha_{sc}\tilde{\alpha}_\delta \Gamma_\times^2 \\
&= A_\delta \cdot 1 - \alpha_{sc}\Gamma_\times \cdot \tilde{\alpha}_\delta \Gamma_\times + \alpha_{sc}\tilde{\alpha}_\delta \Gamma_\times^2 \\
&= A_\delta \qquad \square
\end{aligned}$$

G.5 Instrument power

While the identified object is invariant to κ, the instrument's statistical power is not. The first-stage coefficient is proportional to:

$$\mathrm{Cov}(g_B, Z) \propto \frac{\partial B_B^{sc}}{\partial u_A} = -\alpha_{sc}\, m \left(1 + \frac{N_y \alpha_{sc}}{\Delta}\right) \tag{82}$$

As $\kappa \to \infty$, $m = \partial\mu/\partial u_A \to 0$, and the first stage vanishes. In this limit, the $P_t^{s,\mathrm{USD}}$ channel provides no identifying variation; only cross-book participant reallocation can generate

a non-degenerate first stage.

G.6 Equivalence result

Proposition 7 (Equivalence of microfoundations)**.** *The following two specifications identify the same structural object N_y/A_δ:*

(i) Elastic supply ($P_t^{s,\text{USD}} = 1$) with cross-book participant demand (S_f^{sc}, $\sum_f S_f^{sc} = W$);

(ii) Finite supply elasticity ($\kappa < \infty$) with price-sensitive demand $D_f^{sc}(P_f)$ and no cross-book participants.

In specification (i), the instrument derives power from the budget constraint $\sum_f S_f^{sc} = W$. In specification (ii), the instrument derives power from the supply constraint $S^{sc} = S_0 + \kappa\mu$. Both generate cross-country co-movement in flows through different economic mechanisms but identify the same combination of intermediation cost parameters.

When both channels operate simultaneously (cross-book participants with finite κ), the identified object remains N_y/A_δ, with instrument power reflecting the sum of both channels.

H Microfoundation for cross-book participants

This appendix provides a portfolio-theoretic microfoundation for the cross-book participant demand introduced in Section 3.8. We show that agents who hold fiat balances across multiple currencies optimally reallocate their stablecoin conversions toward currencies experiencing adverse local conditions, generating the cross-country demand reallocation that underpins our identification strategy.

H.1 Agent problem

A representative cross-book participant holds fiat balances $\{N_f^0\}_{f=1}^F$ across F currencies. Each currency f is subject to a stochastic local shock with conversion urgency $\theta_f > 0$. This parameter captures the expected benefit of converting currency f into stablecoins, encompassing expected depreciation, regulatory risk, banking access disruptions, or tightening of capital controls. The residual (unconverted) fiat position in currency f carries variance σ_f^2.

The agent converts c_f units of currency f into USDT at cost y_f per unit (the parity deviation). Total conversion is limited by operational capacity:

$$\sum_f c_f \leq W \tag{83}$$

reflecting exchange limits, KYC processing bandwidth, and settlement constraints.

The agent maximizes expected wealth net of residual risk:

$$\max_{\{c_f \geq 0\}} \quad \sum_f c_f \, \theta_f - \sum_f c_f \, y_f - \frac{\gamma}{2} \sum_f (N_f^0 - c_f)^2 \sigma_f^2 \quad \text{s.t.} \quad \sum_f c_f \leq W \tag{84}$$

The first term captures the expected benefit of conversion: each unit converted avoids the adverse local shock. The second is the premium cost. The third penalizes remaining fiat exposure, weighted by local risk. The parameter $\gamma > 0$ governs risk aversion.

H.2 Optimal allocation

The first-order condition with Lagrange multiplier λ on the budget constraint yields:

$$c_f = N_f^0 + \frac{\theta_f - y_f - \lambda}{\gamma \sigma_f^2} \tag{85}$$

The agent converts more from currency f when conversion urgency θ_f is high, the stablecoin premium y_f is low, local risk σ_f^2 is high, and initial holdings N_f^0 are large.

Imposing $\sum_f c_f = W$ and defining $\Phi \equiv \sum_f (\gamma \sigma_f^2)^{-1}$ and the risk-weighted average $\overline{\theta - y} \equiv \Phi^{-1} \sum_f (\theta_f - y_f)/(\gamma \sigma_f^2)$:

$$S_f^{sc} \equiv c_f = \bar{S}_f^{sc} + \frac{1}{\gamma \sigma_f^2} \left[(\theta_f - y_f) - \overline{\theta - y} \right] \tag{86}$$

where $\bar{S}_f^{sc}$ absorbs the baseline allocation.

H.3 Cross-country transmission

An idiosyncratic shock that raises conversion urgency in country A (θ_A increases) has the following effects:

Own effect:

$$\frac{\partial S_A^{sc}}{\partial \theta_A} = \frac{1}{\gamma \sigma_A^2} \left(1 - \frac{(\gamma \sigma_A^2)^{-1}}{\Phi} \right) > 0 \tag{87}$$

Cross-country effect ($B \neq A$):

$$\frac{\partial S_B^{sc}}{\partial \theta_A} = -\frac{1}{\gamma \sigma_B^2} \cdot \frac{(\gamma \sigma_A^2)^{-1}}{\Phi} < 0 \tag{88}$$

Cross-book participants redirect conversion capacity from country B to country A. The

budget constraint $\sum_f S_f^{sc} = W$ forces the reallocation. The spillover magnitude is proportional to $(\gamma^2\sigma_A^2\sigma_B^2\Phi)^{-1}$: larger when risk aversion is low, both countries' volatilities are low, and fewer currencies compete for conversion capacity.

H.4 Urgency versus premium-chasing

The optimal allocation (86) responds to both urgency (θ_f) and the premium (y_f), but these forces drive reallocation in opposite directions during crisis episodes. When θ_A surges, the urgency channel directs agents to convert more of currency A, drawing resources from other books (S_B^{sc} falls). The resulting rise in y_A simultaneously discourages conversion from A through the $-y_f$ term (a premium-chasing effect that would push agents toward cheaper books). In practice, the urgency channel dominates: expected losses from holding a depreciating or at-risk currency (potentially tens of percentage points during a crisis) far exceed stablecoin premia (typically a few percentage points even under stress).

This has a direct empirical implication. If premium-chasing dominated, a crisis in country A that raises y_A would redirect agents toward books with lower premia, increasing S_B^{sc} and generating a *positive* first stage. The observed *negative* first stage confirms that urgency-driven reallocation is the primary force.

H.5 Reduced form

For the main text, the microfoundation delivers:

$$S_f^{sc} = \bar{S}_f^{sc} + v_f, \qquad \sum_f v_f = 0 \tag{89}$$

where v_f captures idiosyncratic reallocation driven by local conversion urgency. The zero-sum property follows from the finite conversion budget. The main text treats the v_f as given from the stablecoin market's perspective; this appendix establishes that the reallocation is optimal for agents managing fiat risk across multiple currencies.

H.6 Connection to identification

The GIV $Z_f = \sum_{f' \neq f} w_{f'} \hat{\tilde{u}}_{f'}$ captures idiosyncratic urgency shocks $\theta_{f'}$ in other countries (after factor purging absorbs common components). These affect country f's book through the budget constraint: higher $\theta_{f'}$ draws conversion activity from f, lowering S_f^{sc}. The exclusion restriction requires that $\theta_{f'}$ is uncorrelated with country f's local fundamentals. This is plausible because θ_f captures idiosyncratic local conditions rather than global factors, which are removed by the factor-purging step in the instrument construction.