Financial Market Resilience to Intermediary Disruptions: Evidence from Global FX Funding*

Yu An† Amy W. Huber‡

This version: September 2026.

Abstract

How resilient are financial markets to disruptions of the institutions that intermediate them? We answer this question in the global FX funding market using proprietary data that distinguish customers and banks across geography and, for banks, headquarters from physical location. We document strong geographic segmentation among customers and substantial intermediation through cross-border bank operations. These facts motivate an equilibrium model with heterogeneous price responsiveness, which we estimate using granular instrumental variables. We find that cross-border bank operations are more than three times as responsive to funding prices as domestic operations and become important backstops when other suppliers disappear. As a result, dollar funding is resilient to isolated disruptions but becomes increasingly vulnerable as disruptions reach the banks that would otherwise step in. This resilience also limits unilateral financial leverage. Restricting banks headquartered in or operating within the United States generates at most 27% of the price impact from a coordinated restriction on all banks. Financial market resilience therefore depends not simply on who intermediates funding, but on who can replace that intermediation when disruptions occur.

JEL Classifications: F31, G12, G15, G28
Keywords: Financial resilience; Global banking; FX funding demand; CIP deviations; Financial sanctions

*We gratefully acknowledge funding from the Jacobs Levy Equity Management Center for Quantitative Financial Research, and the General Research Support Fund at Johns Hopkins Carey Business School.

†Johns Hopkins Carey Business School. Email: yua@jhu.edu.

‡The Wharton School of the University of Pennsylvania. Email: amyhuber@wharton.upenn.edu.

1 Introduction

Financial markets are networks of specialized institutions. These networks are resilient when, after one institution is disrupted, others can step in and take its place (Allen and Gale 2000; Acemoglu, Ozdaglar, and Tahbaz-Salehi 2015). But who can be replaced depends critically on who else remains in the market. How resilient, then, are financial markets to intermediary disruptions? This question has become increasingly important as financial institutions face threats ranging from cyber and operational failures to financial sanctions and other forms of geoeconomic fragmentation.

We study these questions in the global FX funding market using proprietary data on FX swap positions of 56 customer and bank groups across eight geographic regions. FX funding is synthetic currency funding obtained through FX swaps: investors temporarily obtain one currency against another and reverse the exchange at a predetermined time and exchange rate. The FX funding market is instrumental to cross-border investments (Du and Huber 2024; Huber and Kundu 2026) and is heavily intermediated by banks. A key feature of our data is that we separately observe banks' headquarters and physical location, allowing us to distinguish domestic operations, where the two coincide, from cross-border operations, where they differ. We show that cross-border bank operations account for a substantial share of FX funding and that their funding positions are typically oriented toward the currency of their physical location rather than that of their headquarters. We incorporate these features into an equilibrium model in which investors differ in how strongly they respond to funding prices and in the currency against which they make those adjustments. Using granular instrumental variables (Gabaix and Koijen 2024), we estimate price responsiveness separately for customers, domestic bank operations, cross-border bank operations, and other bank operations. Combined with observed positions and prices, these estimates allow us to recover latent funding demand and reconstruct the equilibrium system, which we use to measure how readily the market can replace lost intermediation.

We find that U.S. dollar (USD) funding is resilient to isolated disruptions, as surviving banks expand their positions to replace lost supply. Cross-border bank operations are particularly important to this resilience: despite their relatively small market share, they respond strongly when other suppliers disappear and become important backstops. The ability of intermediation to migrate across banks limits unilateral financial leverage. Restrictions imposed by any one jurisdiction have only modest effects because banks elsewhere step in, while coordinated restrictions that reach these potential backstops generate substantially larger disruptions.

The data that enable our analysis are daily outstanding FX swap positions in seven

currencies from 2012 to 2024 and are sourced from CLS, the world's largest FX settlement platform. CLS members comprise the major global banks, for which our data identify both headquarters and physical location. Together with geographically identified customers and third-party banks, the resulting panel contains 56 investor groups across eight regions.[1] A second important feature is that we observe the remaining maturity of every outstanding FX swap position each day. We use these data to collapse the entire FX swap book along two dimensions. First, we aggregate across currency pairs to obtain each investor's net funding position in each currency. Second, we use remaining maturity to put contracts of different tenors on a common three-month funding horizon. Contracts that remain outstanding throughout the next three months contribute their full notional, while contracts that expire sooner are weighted by the fraction of the horizon for which they remain outstanding. We thereby obtain, for every investor and currency on every day, a single measure of the quantity of three-month funding.[2]

We document three facts about the geography of global FX funding. First, customers operate in strong home-currency habitats. Across geographic regions, one side of customers' FX funding positions is almost always their home currency. Customers predominantly supply this home currency to obtain funding in foreign currencies, most importantly USD but also EUR and GBP. This pattern suggests that customers enter the FX funding market to manage their home-currency position, making the home currency a natural reference point for their funding decisions.

Second, banks intermediate substantial funding through cross-border operations. Across headquarters jurisdictions, a large share of FX funding activity is conducted outside the bank's home jurisdiction. This is particularly stark for U.S.-headquartered banks, which form the largest banking system in our data and conduct roughly three quarters of their FX funding activity through operations abroad. Cross-border activity concentrates in major host markets, particularly the U.K., creating a considerable separation between where banks are headquartered and where they intermediate FX funding. This separation raises a natural question. Do cross-border operations behave like extensions of their headquarters, or like intermediaries in the markets in which they operate?

Third, cross-border bank operations behave more like local intermediaries than extensions of their headquarters. If funding decisions are centralized within a banking organization, a

[1] The eight regions in our data are U.S., U.K., Japan, the euro area, Switzerland, Canada, Australia, and the rest of the world. The seven currencies are the U.S. dollar (USD), the British pound (GBP), the euro (EUR), the Japanese yen (JPY), the Swiss frank (CHF), the Canadian dollar (CAD), and the Australian dollar (AUD).

[2] FX swap activity is concentrated at relatively short maturities. We use three months rather than a shorter horizon because shorter-horizon funding prices are particularly distorted around quarter-end balance-sheet reporting dates.

foreign operation's funding positions should be concentrated in its headquarters currency. If instead they reflect local funding needs and opportunities, the local currency should dominate. For almost all cross-border operations, we find the latter. The main exceptions are euro-area, Japanese, and U.S. banks operating in the U.K., whose positions remain more concentrated in their headquarters currencies. Together, these facts motivate our model of the FX funding market.

We model the FX funding market as the equilibrium of investor-level funding decisions. Each investor's position reflects an underlying, latent demand for funding and an adjustment to prevailing funding prices, which we measure using deviations from covered interest-rate parity (CIP). The investor adjusts its position as a currency becomes more or less expensive relative to its currency anchor. We allow this price responsiveness to differ across customers, domestic bank operations, cross-border bank operations, and other bank operations. Funding prices adjust so that aggregate demand and supply clear separately in each currency.

We estimate these price responsiveness parameters using granular instrumental variables (GIV). To address the simultaneity between funding quantities and prices, we first remove common components from bank and customer funding movements to isolate idiosyncratic shocks. We then use bank-origin shocks to instrument for funding prices facing customers, and customer-origin shocks to instrument for funding prices facing banks. The identifying assumption is that these idiosyncratic shocks move equilibrium funding prices but are unrelated to the latent funding demand of investors in the other group. We find that cross-border bank operations are more than three times as responsive to funding prices as domestic bank operations, while domestic bank operations, customers, and other banks exhibit similar responsiveness. Combining the estimated responsiveness parameters with observed positions and prices, we recover each investor's latent funding demand and construct the matrix of aggregate price responsiveness across currencies. Together, these objects allow us to resolve equilibrium under counterfactual disruptions to financial intermediation.

We use this estimated system for two counterfactuals. First, we study the resilience of U.S. dollar funding by progressively removing banks that are net suppliers of USD. Removing a bank changes both what the remaining market must accommodate and how responsive the remaining market is. The bank's latent funding demand disappears, changing the former, while its removal changes the composition of investors determining aggregate price responsiveness, thus affecting the latter (Huber and Kundu 2026). We find that the market is highly resilient to isolated disruptions because surviving banks expand their USD supply. As more banks disappear, surviving suppliers take on increasingly large positions, making the system more dependent on those that remain and subsequent disruptions more costly. Cross-border bank operations play a particularly important role in this adjustment. Although some supply

relatively little USD in normal times, their high price responsiveness allows them to become important backstops after other suppliers disappear. Their subsequent removal produces some of the largest jumps in USD funding costs.

Second, we ask how the resilience of the FX funding market constrains the effectiveness of financial sanctions. We prohibit banks headquartered in or operating within the United States ("U.S.-connected banks") from intermediating individual currencies and compare the resulting price effect with a coordinated restriction imposed on all bank suppliers. Despite their large footprints in global FX funding, restricting U.S.-connected banks generates at most 27% of the coordinated effect across the currencies we study. The effect is only 3.4% for EUR and 7.8% even for CAD, which is heavily intermediated by U.S.-connected banks. The contrast is particularly striking for USD. Banks headquartered in or operating within the U.S. provide almost 30% of total bank USD supply, yet eliminating their supply generates only 18% of the price effect achieved under coordinated action. Thus, the same redundancy that makes FX funding resilient to isolated disruptions also limits unilateral financial leverage, with sanctions gaining force only when coordinated across jurisdictions.

This paper is connected to three strands of literature. Our first contribution is to the growing literature on FX derivatives as an important margin of global funding and financial intermediation. FX swaps and related derivatives allow banks and institutional investors to fund and hedge foreign-currency exposures, and recent work shows that activity in these markets reflects important funding and intermediation forces (e.g., Du and Huber 2024; Nenova, Schrimpf, and Shin 2025). In fact, conditions in global FX funding markets can affect equilibrium outcomes in domestic credit markets (e.g., Kubitza, Sigaux, and Vandeweyer 2024; Huber and Kundu 2026). Yet the underlying market structure that determines FX funding conditions remains unclear: who supplies funding, who demands it, and how different participants respond to prices? We characterize this equilibrium structure using proprietary data that are granular in investor geography and observed at high frequency. Our structural estimation recovers heterogeneous price responsiveness and latent funding demand, allowing us to quantify how intermediation is reallocated when part of the market becomes unavailable. This focus on substitution across bank operations within the FX funding market distinguishes our paper from Kloks, Mattille, and Ranaldo (2026) and Khetan (2026), who study how banks substitute away from wholesale dollar funding toward synthetic dollar funding through FX swaps.

We also contribute to the literature on how global banks organize liquidity and transmit shocks across borders. Global banks use internal capital markets and foreign operations to move funding across jurisdictions, which can insulate local activities from some shocks but also transmit shocks across markets (e.g., Schnabl 2012; Cetorelli and Goldberg 2012). We

distinguish the geography of bank ownership from the geography of bank intermediation within the same market. By observing both headquarters and physical location, we show that cross-border bank operations are quantitatively important, often behave more like local intermediaries than extensions of their headquarters, and are substantially more responsive to funding prices than domestic bank operations. This distinction also allows us to ask whether a jurisdiction's reach over globally active banking operations translates into the ability to disrupt funding conditions through financial sanctions. We find that much of the restricted intermediation migrates to banks outside the acting jurisdiction, sharply limiting the impact of unilateral action.

Finally, our paper provides an empirical counterpart to the theoretical literature on the resilience of financial networks. This literature shows that the consequences of a disruption depend critically on the structure of the financial system and on whether surviving institutions can absorb the activity of those that become impaired (e.g., Allen and Gale 2000; Elliott, Golub, and Jackson 2014; Acemoglu, Ozdaglar, and Tahbaz-Salehi 2015; Cabrales, Gottardi, and Vega-Redondo 2017). We quantify this substitution empirically in an important financial market. Our estimated system identifies which intermediaries expand when others become unavailable and how the market's resilience changes as those potential substitutes themselves disappear. We find substantial resilience to isolated disruptions, but this resilience deteriorates as disruptions reach the intermediaries that would otherwise step in. Cross-border bank operations are particularly important along this margin because their high price responsiveness allows them to become important backstops when other suppliers disappear.

The rest of the paper is organized as follows. Section 2 introduces the data and measurement methodology. Section 3 documents three facts about the FX funding market. Section 4 presents the model, which is estimated in Section 5. Section 6 conducts two counterfactual exercises using the estimated model. Section 7 concludes.

2 FX Funding: Concepts and Data

In this section, we define the price and quantity of FX funding and explain how we measure them in the data. Throughout, exchange rates are expressed as number of foreign currencies per USD, and quantities are measured in USD.

2.1 Price of FX Funding

Consider an investor who needs funding in currency n at date t for τ months and uses USD as collateral. At t, she enters an FX swap: sell USD for n at the spot rate $S_{n,t}$ and agree

to buy back USD by delivering n at $t+\tau$ at the forward rate $F_{n,t,t+\tau}$. Let $R_{\$,t,\tau}$ and $R_{n,t,\tau}$ denote gross risk-free rates in USD and n from time t to $t+\tau$. The annualized rate to synthetically borrow 1 unit of currency n using USD is

$$R^{\text{syn}}_{n,t,\tau} \;:=\; R_{\$,t,\tau}\left(\frac{F_{n,t,t+\tau}}{S_{n,t}}\right)^{12/\tau}, \tag{1}$$

where $12/\tau$ annualizes forward points. Covered interest-rate parity (CIP) deviations measure the price of synthetic funding $R^{\text{syn}}_{n,t,\tau}$ relative to borrowing currency n directly at the risk-free rate $R_{n,t,\tau}$, expressed in log differences:

$$P_{n,t,\tau} \;:=\; \ln R^{\text{syn}}_{n,t,\tau} - \ln R_{n,t,\tau}. \tag{2}$$

In a frictionless market satisfying CIP, $P_{n,t,\tau} = 0$, i.e., there is no markup of synthetic funding relative to direct funding. In the data, a negative $P_{n,t,\tau}$ means synthetically funding a foreign currency using USD is cheaper than directly borrowing that foreign currency, which reflects a "USD premium," as documented in post-crisis episodes (Du, Tepper, and Verdelhan, 2018).

Similarly, for any two non-USD currencies n and n', the price of demanding currency n funding and supplying currency n' funding is defined by the triangular arbitrage:

$$P_{n,n',t,\tau} = P_{n,t,\tau} - P_{n',t,\tau}. \tag{3}$$

Data Source and Measurement

We obtain spot and forward FX rates from Bloomberg. To measure risk-free discount rates we use Overnight Indexed Swap (OIS) rates, following Du, Hébert, and Huber (2023). OIS is the appropriate benchmark for CIP because it reflects collateralized overnight funding with minimal bank-credit and scarcity premia; the alternative measure using LIBOR embeds bank credit/liquidity risk (also discontinued around 2021), and measures using Treasury yields can be distorted by specialness and regulatory demand. OIS quotes are available for the USD and six non-USD currencies: Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), euro (EUR), British pound (GBP), and Japanese yen (JPY).

A contribution of this paper is to construct a continuous OIS-based series of CIP deviations that spans the CHF and EUR reference-rate transitions. Because these transitions induce level shifts, many studies restrict CHF and EUR samples to single-reference subperiods. We instead splice the pre- and post-transition segments using the overlap window to align levels at each tenor, yielding a single continuous series.

Our sample runs from September 2012 through December 2024, with the start date dictated by the availability of the quantity data described below. We study the most liquid tenors, including 1-month (1M), 3-month (3M), 6-month (6M), and 1-year (1Y). Appendix Figure A1 plots $P_{n,t,3\mathrm{M}}$, the 3M price of demand funding in currency n by supplying USD funding for each non-USD currency n, as defined in (2). We observe pronounced comovement across currencies, consistent with tightly connected FX funding markets exposed to common supply- and demand-side shocks. Results for other tenors are similar and omitted for brevity.

2.2 Quantity of FX Funding

Our objective is to recover the total quantity relevant for pricing at a given funding horizon. Positions at longer maturities continue to occupy balance-sheet capacity within shorter horizons as they roll down. For example, even if an investor has no outstanding 1-month position today, his balance sheet can be occupied for the next 1-month horizon from a recently entered 3-month USD borrowing. Accordingly, our quantity measure over funding horizon τ aggregates outstanding contracts across all remaining maturities, weighting each position by the share of its life that overlaps $[t, t+\tau]$. This roll-down mapping aligns quantities with the horizon-specific price $P_{n,t,\tau}$ and avoids undercounting by looking only at contracts with tenor τ.

For investor group i demanding funding in currency n on date t, the horizon-τ quantity measure is

$$Q^D_{i,n,t,\tau} := \sum_{m\in\{1\mathrm{W},1\mathrm{M},3\mathrm{M},6\mathrm{M},1\mathrm{Y},\mathrm{Longer}\}} O^D_{i,n,t,m}\,\mu^D_{i,n,t,m}(\tau), \quad \mu^D_{i,n,t,m}(\tau) := \frac{\min\{\tau,\, M^D_{i,n,t,m}\}}{\tau} \in [0,1], \tag{4}$$

where m indexes the maturity buckets, $O^D_{i,n,t,m}$ is the end-of-day outstanding total notional of currency n demanded by investor i on day t in bucket m, and $M^D_{i,n,t,m}$ is the corresponding outstanding-amount-weighted remaining maturity. The weight $\mu^D_{i,n,t,m}(\tau)$ equals 1 when the remaining maturity exceeds τ and scales linearly otherwise. Intuitively, a 1-month contract contributes 1/12 to the 1-year quantity, whereas a 1-year contract contributes fully (weight 1) to the 1-month quantity. We construct the supply-side quantity analogously, $Q^S_{i,n,t,\tau} = \sum_m O^S_{i,n,t,m}\,\mu^S_{i,n,t,m}(\tau)$; the net funding quantity demanded by investor i is $Q_{i,n,t,\tau} = Q^D_{i,n,t,\tau} - Q^S_{i,n,t,\tau}$.

Data Source and Measurement

Our quantity data come from the CLS Group (CLS), which settles FX trades for 72 settlement members, primarily large multinational banks.[3] As the largest single source of FX execution data, CLS covers over 50% of global FX volumes.

We obtain from CLS daily aggregates of outstanding FX swap and forward positions. The data are organized by date, investor category, physical location, currency pair, and maturity bucket. For each cell, we observe the total outstanding notional on the buy and sell sides, aggregated across counterparties, and the outstanding-notional-weighted average remaining maturity. Remaining maturity is essential for converting contracts with different maturities into a common-horizon measure of FX funding.[4]

CLS distinguishes banks from nonbank customers and reports the physical location of their funding activity in eight regions: the United States (US), euro area (EZ), United Kingdom (UK), Japan (JP), Australia (AU), Switzerland (CH), Canada (CA), and rest of world (ROW). For customers, only physical location is observed, so each region constitutes a separate customer group. For CLS-member banks, which settle directly through CLS, we observe both headquarters and physical location and treat each headquarters-by-location cell as a separate bank operation. Among the seven individually identified regions, activity conducted in a bank's headquarters region is classified as *domestic bank operations*, while activity conducted outside its headquarters region is classified as *cross-border bank operations*. For banks that do not settle directly through CLS, headquarters is not observed; we refer to these as third-party banks and group them by physical location. Third-party banks and bank operations involving ROW constitute *other bank operations*. The resulting active sample contains 56 investor groups: 8 customer regions, 7 domestic bank operations, 24 cross-border bank operations, and 17 other bank operations.

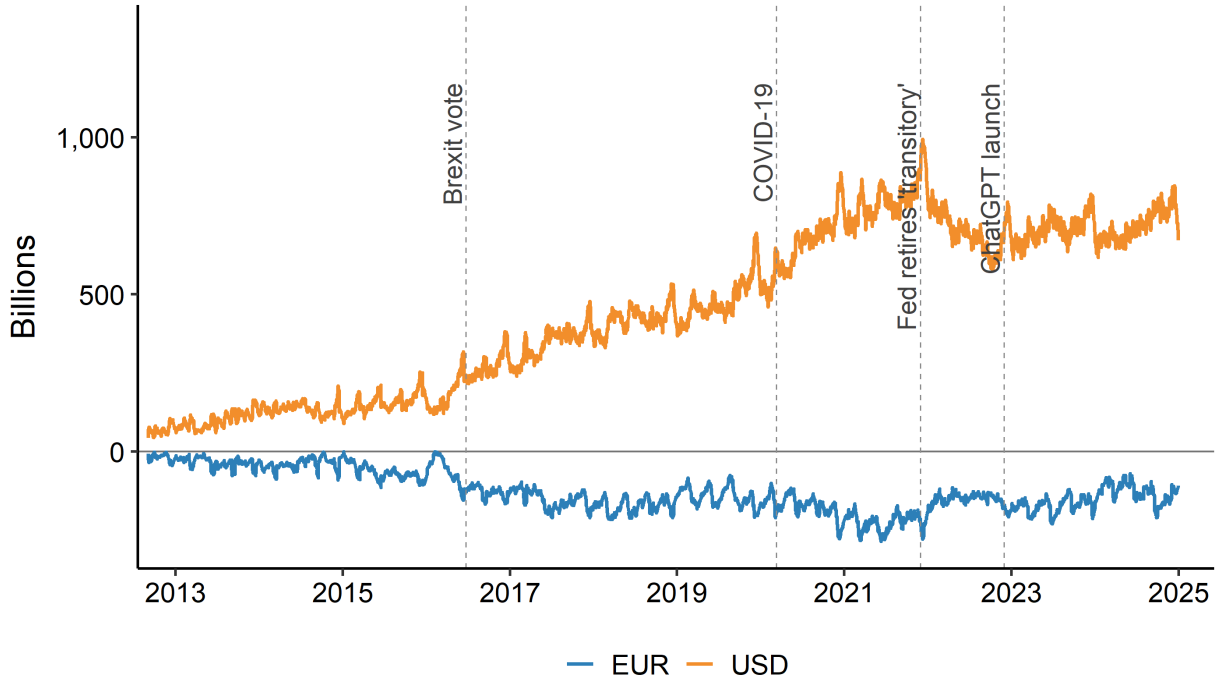
The raw CLS panel covers 17 currencies from September 2012 through December 2024. To match the availability of our funding-price benchmarks, we restrict the analysis to USD, AUD, CAD, CHF, EUR, GBP, and JPY, the seven currencies for which we have OIS curves.[5] We convert currency-pair positions into currency-level funding demand and supply. A posi-

[3] A list of settlement members is available at https://www.cls-group.com/communities/settlement-members/.

[4] CLS defines remaining maturity as the number of calendar days from the current date to settlement of the far leg, which typically occurs two business days after the contractual far date. The maturity buckets are tomorrow-next and spot-next, with remaining maturity of at most 3 days; 1 week, 4–9 days; 1 month, 10–33 days; 3 months, 34–94 days; 6 months, 95–185 days; 1 year, 186–368 days; and longer than 1 year, at least 369 days.

[5] The remaining CLS currencies are Danish krone (DKK), Hong Kong dollar (HKD), Israeli new shekel (ILS), Korean won (KRW), Mexican peso (MXN), New Zealand dollar (NZD), Norwegian krone (NOK), Swedish krona (SEK), Singapore dollar (SGD), and South African rand (ZAR).

Figure 1: FX funding quantity in USD and EUR



Notes: This figure plots the time series of aggregate FX funding quantities in USD and EUR at the 3-month horizon, measured as net customer demand for funding from banks in each currency and expressed in billions of USD. Positive values indicate that customers are net borrowers of the currency (net funding demand), while negative values indicate that they are net lenders (net funding supply).

tion that borrows currency n against currency n' contributes positively to funding demand in n and negatively to funding demand in n'. For example, borrowing EUR against JPY contributes to EUR funding demand and JPY funding supply. We express all quantities in USD using contemporaneous spot exchange rates, consistent with our USD-numeraire convention.

Figure 1 shows the evolution of aggregate FX funding quantities in USD and EUR. Over 2013-2024, net investor demand for USD funding rises almost monotonically from about $100 billion to more than $750 billion, with noticeable accelerations around the Brexit vote and the COVID-19 shock. The series dips temporarily after the U.S. inflation was declared non-transitory before increasing again alongside the recent AI-related rally. By contrast, investors are net suppliers of EUR funding throughout the sample: EUR quantities are negative, becoming more negative after COVID-19 and then partly reversing after the inflation was declared non-transitory. The event markers line up with several turning points in both series, suggesting that major political and macroeconomic developments reshape the demand for dollar and euro funding.

3 The Geography of FX Funding

Our data reveal a pronounced geographic organization of the FX funding market. Customers operate in strong home-currency habitats, while CLS-member banks intermediate substantial funding through operations located outside their headquarters jurisdictions. At the same time, cross-border bank operations do not simply inherit the currency orientation of their parent banking organizations: their funding portfolios are typically more closely oriented toward the currency of the market in which they operate. We document these three facts before introducing the model.

Fact 1: Customers operate in strong home-currency habitats

We begin with customers. For each customer region, we summarize its average funding portfolio by dividing the time-series average position in each currency by the region's average one-sided gross funding, where one-sided gross funding equals one half of the sum of absolute currency positions. The home-currency share is the absolute value of this normalized position for the region's home currency. A value of 100 percent means that the home currency constitutes one entire side of the customer's average net funding portfolio.

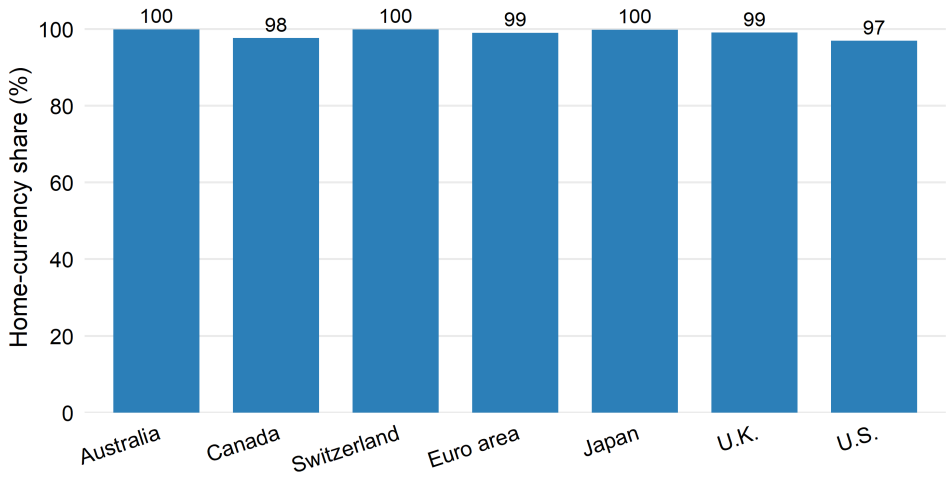
Figure 2a shows striking home-currency concentration. Across all customer regions with an observed home currency, the home-currency share is close to 100 percent. Thus, despite participating in a global market and holding positions in multiple currencies, customers' funding portfolios exhibit an almost complete home-currency orientation. This preferred habitat is common across regions rather than being a feature of one or two particularly large customer groups.

The home-currency concentration does not imply that customer funding consists exclusively of home-currency–USD positions. Figure 2b shows the full currency composition of each region's average funding portfolio using the same normalization as Figure 2a.

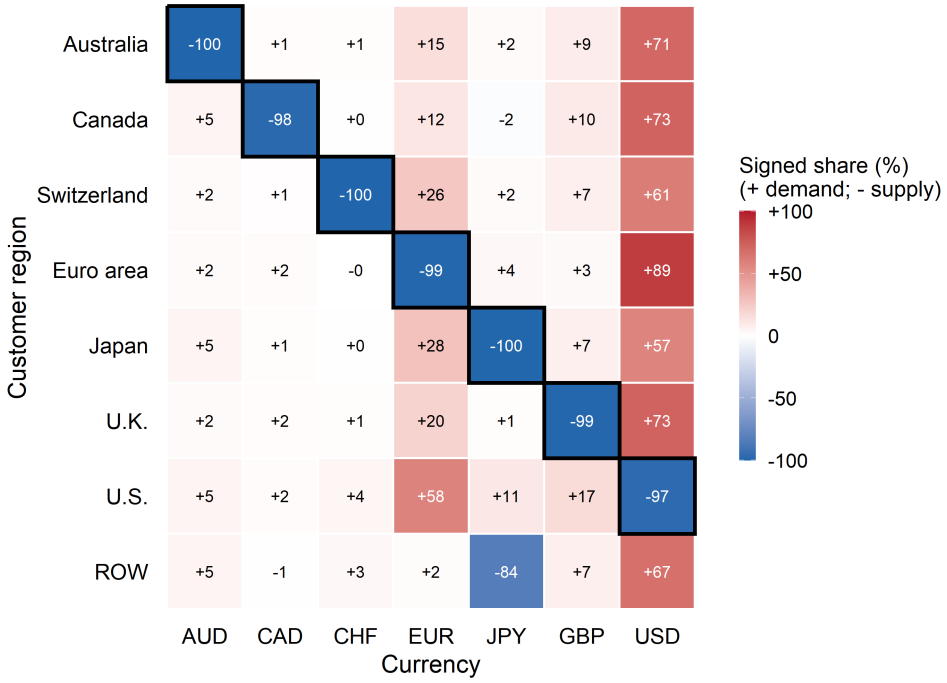
The heatmap reveals a common structure. The home currency forms a dominant leg of the funding position in every region. USD frequently accounts for a large part of the opposite side, consistent with its central role in the FX swap market, but the counterposition is not exclusively in USD: other currencies make economically meaningful contributions, and their importance differs across regions. The salient common feature is therefore not simply USD dominance, but the distinct home-currency habitat from which each customer region enters the multicurrency funding market.

Figure 2: Customer home-currency habitats

(a) Home-currency concentration



(b) Currency composition of customer funding positions



Notes: Figure 2a reports the home-currency share for each of the seven customer regions with a unique geographic home currency; ROW is omitted because it has no unique home currency. The home-currency share is the absolute value of the region's average home-currency position divided by its average one-sided gross funding. Figure 2b reports each region's average signed position in each currency using the same normalization. Black outlines identify geographic home currencies, so the absolute value of each outlined cell corresponds to the bar in Figure 2a. One-sided gross funding equals one half of the sum of absolute currency positions. Positive entries denote net demand and negative entries denote net supply. Positions are three-month FX-funding positions over the full sample.

Fact 2: Banks operate substantial funding books outside their headquarters jurisdictions

The geographic organization of banks differs sharply from that of customers. For each CLS-member bank operation, we observe both the headquarters jurisdiction of the banking organization and the physical location in which the position is booked. This distinction allows us to separate where a bank is headquartered from where it operates its FX funding book.

Figure 3a shows that cross-border operations account for a substantial share of CLS-member bank funding across headquarters jurisdictions, although the relative importance of domestic and cross-border activity varies across banking systems. Consolidating all positions at the headquarters level would therefore obscure an important part of the market through which global FX funding is intermediated.

Figure 3b focuses on cross-border operations and organizes them by physical location. The U.K. stands out as the largest booking location by a wide margin, while the euro area is also an important center of cross-border bank activity. Both locations host substantial funding books operated by banks headquartered in several different jurisdictions, consistent with London and major euro-area financial centers serving as hubs for global bank intermediation. More generally, headquarters and physical location capture distinct dimensions of bank activity: headquarters identifies the banking organization that owns the operation, while physical location identifies where its funding book is operated.

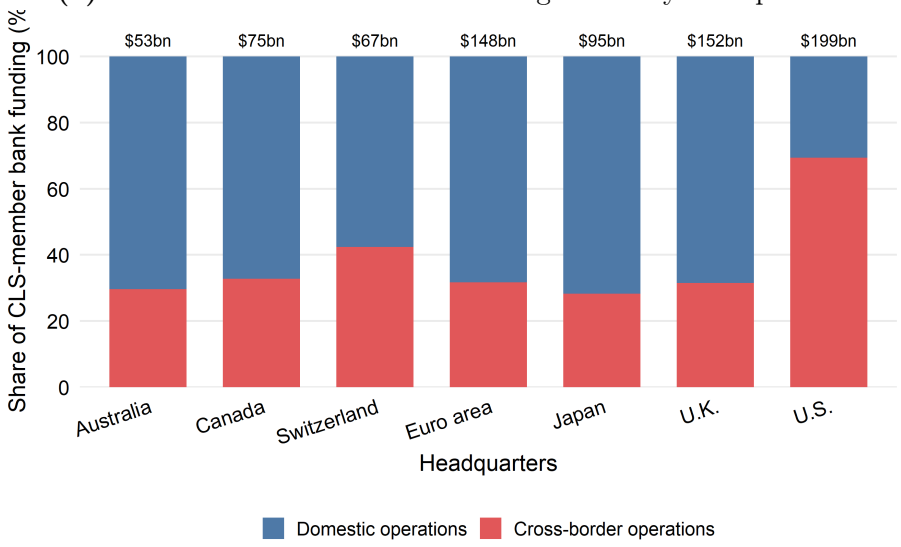
Fact 3: Cross-border bank operations typically have a local currency anchor

The distinction between headquarters and physical location raises a further question. When a CLS-member bank operates outside its headquarters jurisdiction, which currency best organizes its funding portfolio? An integrated-bank view suggests the headquarters currency as a natural candidate: cross-border operations belong to a consolidated banking organization whose funding costs, capital allocation, and ultimate profits may be evaluated at headquarters. A local-intermediation view instead suggests the currency of the operation's physical location, reflecting local customers, funding conditions, and balance-sheet management.

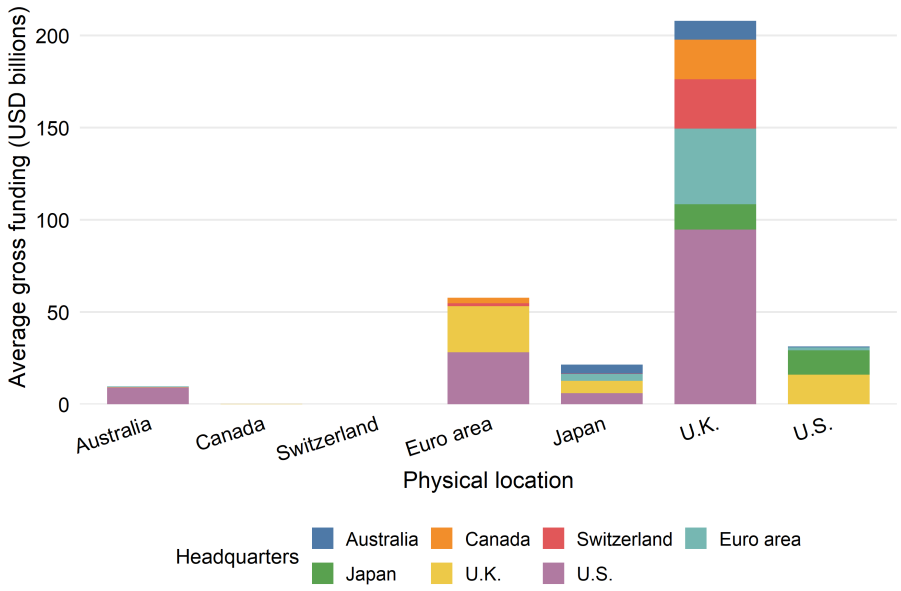
We let the positions speak to this distinction. For each cross-border CLS-member bank operation whose headquarters and physical location are among the seven individually identified regions, we compare the home-currency shares associated with the headquarters and physical-location currencies. For either candidate, the home-currency share is the sum over time of the operation's absolute position in that currency divided by the corresponding sum of one-sided gross funding. Figure 4a plots the two shares against each other. Operations

Figure 3: The geographic organization of CLS-member bank funding

(a) Domestic and cross-border funding shares by headquarters



(b) Headquarters composition of cross-border operations by physical location



Notes: Figure 3a reports, for each headquarters jurisdiction, the shares of its CLS-member bank funding book operated domestically and across borders; the two shares sum to 100 percent. Labels above the bars report the total size of the combined domestic and cross-border funding book in USD billions. Figure 3b reports only cross-border CLS-member bank operations by physical location, with each bar decomposed by headquarters. Both figures include operations whose headquarters and physical locations are among the seven individually identified regions; operations involving ROW and third-party banks are omitted. One-sided gross funding equals one half of the sum of absolute currency positions.

above the 45-degree line are more strongly oriented toward their physical-location currency, while those below it are more strongly oriented toward their headquarters currency. Marker size reflects the scale of the operation's funding position.

Figure 4a shows that the physical-location currency provides the stronger orientation for most cross-border bank operations, often by a substantial margin. These operations therefore generally look more like participants in the funding market in which they operate than like mechanically replicated extensions of their headquarters funding book. The principal exceptions are several economically important operations located in the U.K., for which the headquarters currency remains stronger, consistent with London's distinctive role as an international platform on which banks headquartered in many jurisdictions intermediate global funding.

We translate this evidence into a currency anchor for each investor. A domestic CLS-member bank operation is assigned the common currency of its headquarters and physical location. A cross-border CLS-member bank operation whose headquarters and physical location are both individually identified is assigned whichever of those two currencies has the larger home-currency share; ties are assigned to the physical-location currency. Figure 4b summarizes these assignments across the headquarters-by-physical-location matrix.[6]

The comparison clarifies both the strength and the limits of the local result. Physical location dominates for many cross-border bank operations rather than winning mechanically by small differences in portfolio shares. At the same time, the economically important exceptions demonstrate why neither headquarters nor physical location can be imposed universally as the relevant currency reference point.

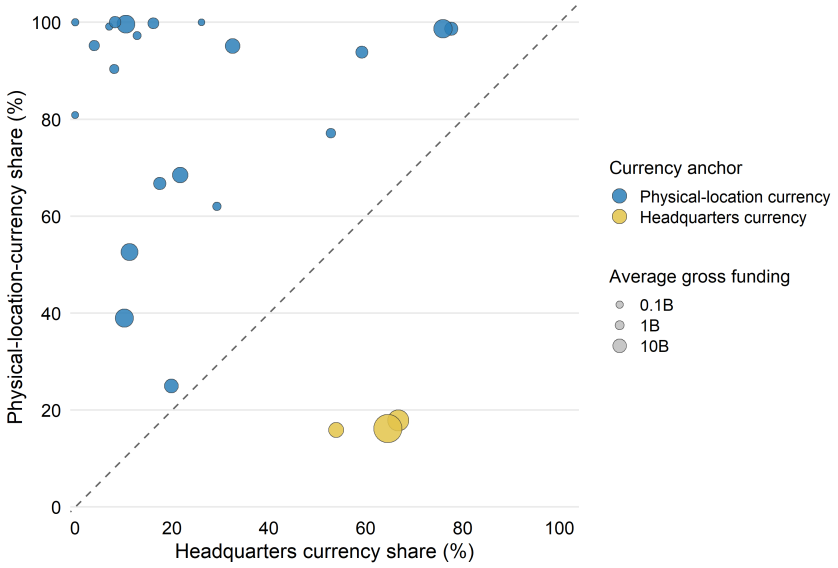
We interpret these patterns as evidence about the organization of observed funding portfolios, not as direct evidence about the internal numeraire used by a bank's treasury or trading desk. Headquarters may continue to determine consolidated P&L, internal transfer prices, or capital allocation even when a cross-border operation's FX funding book is primarily organized around its local currency.

Together, these facts reveal a market organized simultaneously by geography and organizational structure. Customers enter the market from sharply defined home-currency habitats. CLS-member banks connect these geographically segmented investors through a global network of operations. Within that network, the currency anchor of an operation

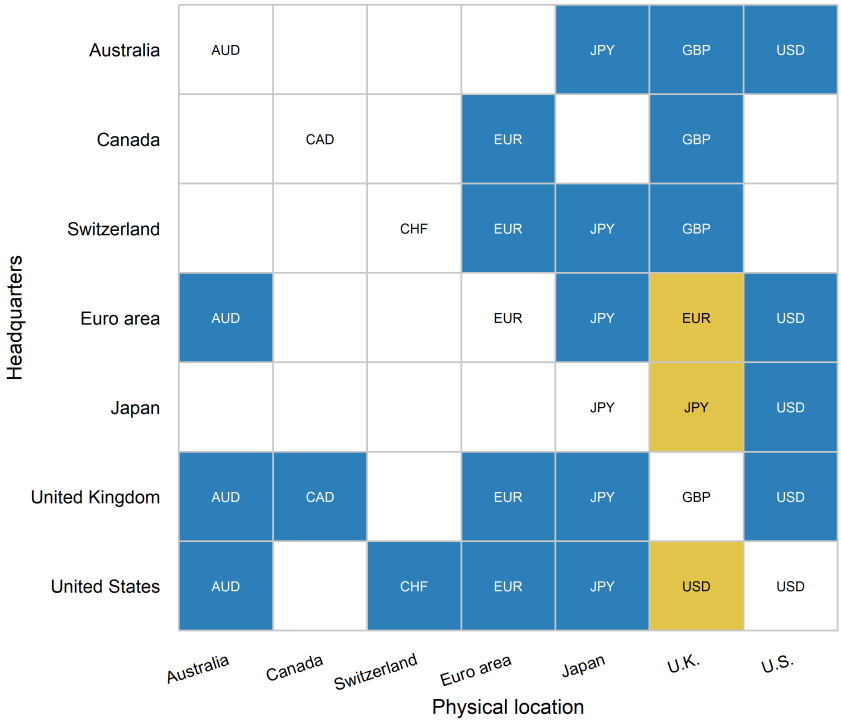
[6]Customer regions with an individually identified home currency are assigned that currency, while the ROW customer is assigned the currency with the largest home-currency share. For a CLS-member bank operation involving ROW on only one side, we assign the currency associated with the individually identified headquarters or physical location; when both headquarters and physical location are ROW, we assign the currency with the largest home-currency share. For a third-party bank, whose headquarters is unobserved, we assign the physical-location currency when it is individually identified and otherwise assign the currency with the largest home-currency share.

Figure 4: Currency anchors of CLS-member bank operations

(a) Physical-location- versus headquarters-currency orientation



(b) Currency anchors by headquarters and physical location



Notes: Figure 4a compares the headquarters-currency and physical-location-currency shares for cross-border CLS-member bank operations. Points above the 45-degree line are more strongly oriented toward the physical-location currency and shown in blue; points below the line are more strongly oriented toward the headquarters currency and shown in yellow. Marker size reflects average one-sided gross funding. Figure 4b reports the assigned currency anchors across headquarters-by-physical-location cells. Both figures restrict attention to CLS-member bank operations whose headquarters and physical locations are among the seven individually identified regions; operations involving ROW and third-party banks are omitted.

is typically associated with where it operates, although important headquarters-oriented exceptions remain.

4 A Model of Global FX Funding

The empirical facts in Section 3 point to a simple organization of the market. Customer funding positions are anchored in home currencies, while banks connect these regional habitats through domestic and cross-border operations. Most cross-border bank operations are more strongly oriented toward the currency of their physical location, although some retain a headquarters orientation. We capture these features by treating each customer region and bank operation as a separate investor with its own currency anchor, while allowing price responsiveness to differ across the four investor groups documented in the data. The model then separates latent demand from the response of positions to relative funding prices.

4.1 Positions and Currency Anchors

Let currencies be indexed by $n = 0, \ldots, N$, with USD indexed by 0. We normalize $P_{0,t} = 0$, where $P_{n,t}$ is the price of obtaining currency n funding against USD. The price of obtaining currency n funding against currency h is therefore $P_{n,t} - P_{h,t}$.

Investors are indexed by $i = 1, \ldots, I$. An investor is either a customer region or a bank operation. For CLS-member banks, each headquarters-by-physical-location cell constitutes a separate operation; third-party banks are grouped by physical location. Treating these operations separately reflects the evidence in Figure 3 that banks operate substantial funding books outside their headquarters jurisdictions. Let $h(i)$ denote investor i's currency anchor, assigned from full-sample positions as described in Section 3 and held fixed throughout the analysis. For customers, these assignments reflect the home-currency habitats in Figure 2; for CLS-member bank operations, they reflect the physical-location-versus-headquarters evidence summarized in Figure 4.

Let $Q_{i,n,t}$ denote investor i's signed net demand for funding in currency n at time t. Positive values indicate demand for currency n, while negative values indicate supply. Because an FX funding position pairs demand in one currency with supply in another, each investor's positions add up to zero:

$$\sum_{n=0}^{N} Q_{i,n,t} = 0. \tag{5}$$

Funding positions differ sharply in scale across customer regions and bank operations. Let $S_{i,t-1} > 0$ denote the size of investor i at time $t-1$, and write its position per unit of

lagged size as

$$Q_{i,n,t} = S_{i,t-1} q_{i,n,t}. \tag{6}$$

This normalization makes price responsiveness comparable across investors of different sizes. Define investor i's set of non-anchor currencies as

$$\mathcal{N}_i := \{0, \ldots, N\} \setminus \{h(i)\}. \tag{7}$$

4.2 Latent Demand, Price Adjustment, and Equilibrium

Let $g(i) \in \{C, D, X, O\}$ denote investor i's group. We distinguish customers, domestic bank operations, cross-border bank operations, and other bank operations. The last group comprises third-party banks and CLS-member bank operations involving ROW. We impose common price responsiveness within each group:

$$\theta_{g(i)} = \begin{cases} \theta_C, & i \text{ is a customer region,} \\ \theta_D, & i \text{ is a domestic bank operation,} \\ \theta_X, & i \text{ is a cross-border bank operation,} \\ \theta_O, & i \text{ is an other bank operation.} \end{cases} \tag{8}$$

This grouping is deliberately parsimonious. It distinguishes customers from banks and, within banks, separates domestic and cross-border operations as motivated by Figure 3. An operation's group determines its price responsiveness but not its currency anchor. In particular, a cross-border bank operation may be anchored to either its physical-location currency or its headquarters currency while retaining the common cross-border price responsiveness θ_X.

For each $n \in \mathcal{N}_i$, investor i's normalized funding position satisfies

$$q_{i,n,t} = d_{i,n,t} - \theta_{g(i)} \left(P_{n,t} - P_{h(i),t} \right), \qquad n \in \mathcal{N}_i. \tag{9}$$

Here, $d_{i,n,t}$ is investor i's latent demand for currency n funding per unit of size. For customers, latent demand reflects investment, hedging, and operating needs. For banks, it reflects funding advantages, balance-sheet conditions, and the scale of intermediation they seek to provide. More generally, $d_{i,n,t}$ absorbs all time-varying determinants of positions other than the contemporaneous bilateral funding price.

The parameter $\theta_{g(i)} > 0$ measures the price responsiveness of normalized positions. When currency n becomes more expensive relative to the investor's currency anchor, the investor demands less of it or supplies more of it. Because the dependent variable is a signed position

normalized by investor size and the price variable is a bilateral funding-price spread, $\theta_{g(i)}$ is a demand slope rather than a conventional percentage-change elasticity. Equation (9) is therefore a simple linear demand schedule in which each non-anchor position responds to its own bilateral funding price relative to the currency anchor.

The anchor-currency position is determined by adding up:

$$q_{i,h(i),t} = -\sum_{n \in \mathcal{N}_i} q_{i,n,t}. \tag{10}$$

We define latent demand for the anchor currency analogously, so latent demands also add up to zero. The model therefore has N independent position choices for each investor.

The model does not impose a fixed foreign-currency basket. Each $d_{i,n,t}$ can vary separately across currencies and over time. This flexibility matches Figure 2b: the currency anchor forms one dominant leg of the customer position, while the opposite leg spans several currencies and differs across regions. The model thus combines a persistent currency anchor with a flexible counterbasket. Investors remain heterogeneous through their currency anchors, sizes, and latent demands, even when they share common price responsiveness within an investor group.

At each time t, latent demands and lagged investor sizes are taken as given, and relative funding prices adjust to clear each currency market:

$$\sum_{i=1}^{I} Q_{i,n,t} = \sum_{i=1}^{I} S_{i,t-1} q_{i,n,t} = 0, \qquad n = 0, \ldots, N. \tag{11}$$

Because every investor's positions add up to zero, only N of these equations are independent. Together with the normalization $P_{0,t} = 0$, they determine equilibrium relative funding prices.

5 Estimation and Equilibrium Structure

We now take the model to the data. We estimate the four group-specific price-responsiveness parameters in equation (9), use them to recover investor-level latent demand, and aggregate the estimated demand system into a market-level equilibrium. The key empirical challenge is simultaneity: shocks to latent demand move both investors' positions and equilibrium funding prices, so a regression of position changes on price changes does not isolate the price-induced adjustment in the demand equation. We address this challenge using granular instruments constructed from idiosyncratic shocks on the opposite side of the market.

5.1 Empirical Setup

Throughout this section, t indexes calendar months. We estimate the model using three-month positions and funding prices for USD, AUD, CAD, CHF, EUR, GBP, and JPY. In each month, we average positions and prices over the same first five dates on which the complete seven-currency price vector is observed. We then estimate the model using monthly changes in positions normalized by lagged investor size.

We proxy investor size by its gross funding position and normalize current positions by lagged size:

$$S_{i,t} := \sum_{n=0}^{N} |Q_{i,n,t}|, \qquad q_{i,n,t} := \frac{Q_{i,n,t}}{S_{i,t-1}}. \tag{12}$$

For each non-anchor currency, the monthly estimating equation is therefore

$$\Delta q_{i,n,t} = \Delta d_{i,n,t} - \theta_{g(i)} \Delta \left(P_{n,t} - P_{h(i),t} \right), \qquad n \in \mathcal{N}_i. \tag{13}$$

We use the currency-anchor assignments documented in Section 3 and hold them fixed throughout estimation. Customer regions enter separately and share the common price responsiveness θ_C. For bank operations, we aggregate raw positions and size within each bank-operation group and assigned currency anchor before estimating θ_D, θ_X, and θ_O. Thus, an operation's group determines its price responsiveness, while its currency anchor determines the relevant bilateral funding price.

The bilateral funding price in equation (13) is endogenous because it responds to the same latent-demand shifts contained in the error term. We address this problem using granular instruments (Gabaix and Koijen, 2024). For each currency, we remove common components from position changes on one side of the market and aggregate the remaining investor-specific shocks using lagged size weights. Customer equations are instrumented with bank-origin shocks, while bank equations are instrumented with customer-origin shocks. For a position in currency n relative to currency anchor h, the instrument is the difference between the source-side granular shocks to n and h. We control for the source-side common factors and cross-fit the pair-specific first-stage weights. Identification requires that, conditional on these common factors, opposite-side idiosyncratic position shocks affect the target investors' positions only through bilateral funding prices. Appendix A provides the full construction and inference details.

Table 1: Granular-IV Estimates of FX Funding Price Responsiveness

	Customers	Domestic bank operations	Cross-border bank operations	Other bank operations
Price responsiveness	0.37*	0.23**	0.89***	0.28**
	(0.22)	(0.11)	(0.32)	(0.12)
Observations	6,948	6,132	4,548	6,132
First-stage F	11.19	11.42	10.70	11.42
	Wald statistic		p-value	
Equal bank price responsiveness	5.23		0.07	

Notes: The table reports estimates of the price-responsiveness parameter θ_g in equation (13). Newey–West standard errors are reported in parentheses. The bandwidth is selected using the rule of thumb $L_T = \lfloor 4(T/100)^{2/9} \rfloor$, which yields four monthly lags in our sample. The first-stage F-statistics are HAC Wald statistics for the cross-fitted scalar instrument and use the same bandwidth. The lower portion reports a HAC Wald test of equality across the price-responsiveness coefficients of domestic, cross-border, and other bank operations, with the p-value reported alongside the test statistic. ***, **, and * denote statistical significance at the 1, 5, and 10 percent levels, respectively.

5.2 Granular-IV Estimates of Price Responsiveness

Table 1 reports the granular-IV estimates of price responsiveness for the four investor groups. All four estimates are positive and statistically significant at the 10 percent level, and the first-stage F-statistics range from 10.70 to 11.42.

The estimated price responsiveness is 0.37 for customers. This means that a one-basis-point increase in a currency's funding price relative to the investor's currency anchor reduces demand for that currency, or increases its supply, by 0.37 percent of lagged investor size. Among bank operations, the corresponding responses are 0.23 percent for domestic operations, 0.89 percent for cross-border operations, and 0.28 percent for other operations. Thus, a 10-basis-point increase induces position adjustments equal to approximately 3.7, 2.3, 8.9, and 2.8 percent of lagged size, respectively. Cross-border bank operations are therefore much more responsive to relative funding prices: their estimated coefficient is approximately four times that of domestic bank operations and more than three times that of other bank operations.

The equality test rejects equality across the three bank-operation coefficients at the 10 percent level, with a Wald statistic of 5.23 and a p-value of 0.07. Together with the descriptive evidence in Section 3, the estimates indicate that cross-border bank operations not only account for substantial global funding activity but also adjust their positions more strongly in response to relative funding prices.

5.3 Aggregate Latent Demand and Market Responsiveness

The estimated price responsiveness allows us to invert the demand equation and separate latent demand from the price-induced component of observed positions. For each active investor and non-anchor currency, we recover

$$\widehat{d}_{i,n,t} = q_{i,n,t} + \widehat{\theta}_{g(i)} \left(P_{n,t} - P_{h(i),t} \right), \qquad n \in \mathcal{N}_i. \tag{14}$$

The anchor-currency component is determined by adding up and therefore does not provide an additional independent observation.

The recovered latent demands and estimated price responsiveness allow us to aggregate the investor-level demand system into a market-level equilibrium. Let $\mathbf{P}_t$ collect the $N+1$ currency funding prices, and let $\widehat{\mathbf{d}}_{i,t}$ collect investor i's recovered latent demand per unit of size for its N non-anchor currencies. For each possible currency anchor h, define the $N \times (N+1)$ anchor-incidence matrix $\mathbf{H}_h$. The row associated with non-anchor currency n has a $+1$ in the column for currency n, a -1 in the column for anchor currency h, and zeros elsewhere. Thus, $\mathbf{H}_h \mathbf{P}_t$ collects the bilateral prices $P_{n,t} - P_{h,t}$.

As shown in Appendix A.4, each investor's full currency-position vector equals its size-scaled latent demand minus the position adjustment induced by relative prices. Summing across investors and imposing market clearing gives

$$\mathbf{A}_t \mathbf{P}_t = \mathbf{b}_t, \qquad \mathbf{b}_t := \sum_i S_{i,t-1} \mathbf{H}_{h(i)}^\top \widehat{\mathbf{d}}_{i,t}, \qquad \mathbf{A}_t := \sum_i S_{i,t-1} \widehat{\theta}_{g(i)} \mathbf{H}_{h(i)}^\top \mathbf{H}_{h(i)}. \tag{15}$$

The vector $\mathbf{b}_t$ is aggregate latent demand. Each component of an investor's non-anchor latent demand is assigned to the corresponding currency and offset by an equal position in its currency anchor. Multiplication by $S_{i,t-1}$ converts latent demand per unit of size into position units. The matrix $\mathbf{A}_t$ is the aggregate position-response matrix. Investor i contributes its size-weighted price responsiveness, $S_{i,t-1}\widehat{\theta}_{g(i)}$, to each bilateral margin connecting its currency anchor to one of its non-anchor currencies. Summing these contributions gives the market-wide response of positions to relative funding prices. Equation (15) therefore states that equilibrium relative prices generate a position response, $-\mathbf{A}_t \mathbf{P}_t$, that offsets aggregate latent demand. Because only relative prices matter, we normalize the USD funding price to zero.

Table 2 shows that the strongest aggregate responses involve EUR and USD. These currencies have the largest diagonal entries, and EUR–USD is the strongest individual bilateral margin. The matrix is nevertheless dense rather than organized exclusively around USD, reflecting the heterogeneous currency anchors documented in Section 3. Its cross-currency structure is also stable over time: a fixed matrix shape scaled by a single monthly factor

Table 2: Aggregate Position-Response Matrix Across Currencies

	USD	AUD	CAD	CHF	EUR	GBP	JPY
USD	19.5	−2.7	−2.6	−2.7	−4.6	−3.6	−3.2
AUD		10.7	−0.8	−0.9	−2.9	−1.9	−1.4
CAD			10.1	−0.8	−2.8	−1.8	−1.3
CHF				10.7	−2.9	−1.9	−1.4
EUR					20.4	−3.8	−3.4
GBP						15.4	−2.4
JPY							13.1

Notes: The table reports the upper triangle of the time-series average of $\mathbf{A}_t$, normalized so that its trace equals 100. A diagonal entry sums the aggregate position responses on all bilateral margins involving that currency. Off-diagonal entries are negative; their absolute values measure aggregate responsiveness on the corresponding currency pair. The lower triangle is omitted because $\mathbf{A}_t$ is symmetric.

explains about 88 percent of the variation in $\mathbf{A}_t$. We nevertheless use the month-specific aggregate position-response matrix in all subsequent counterfactuals.

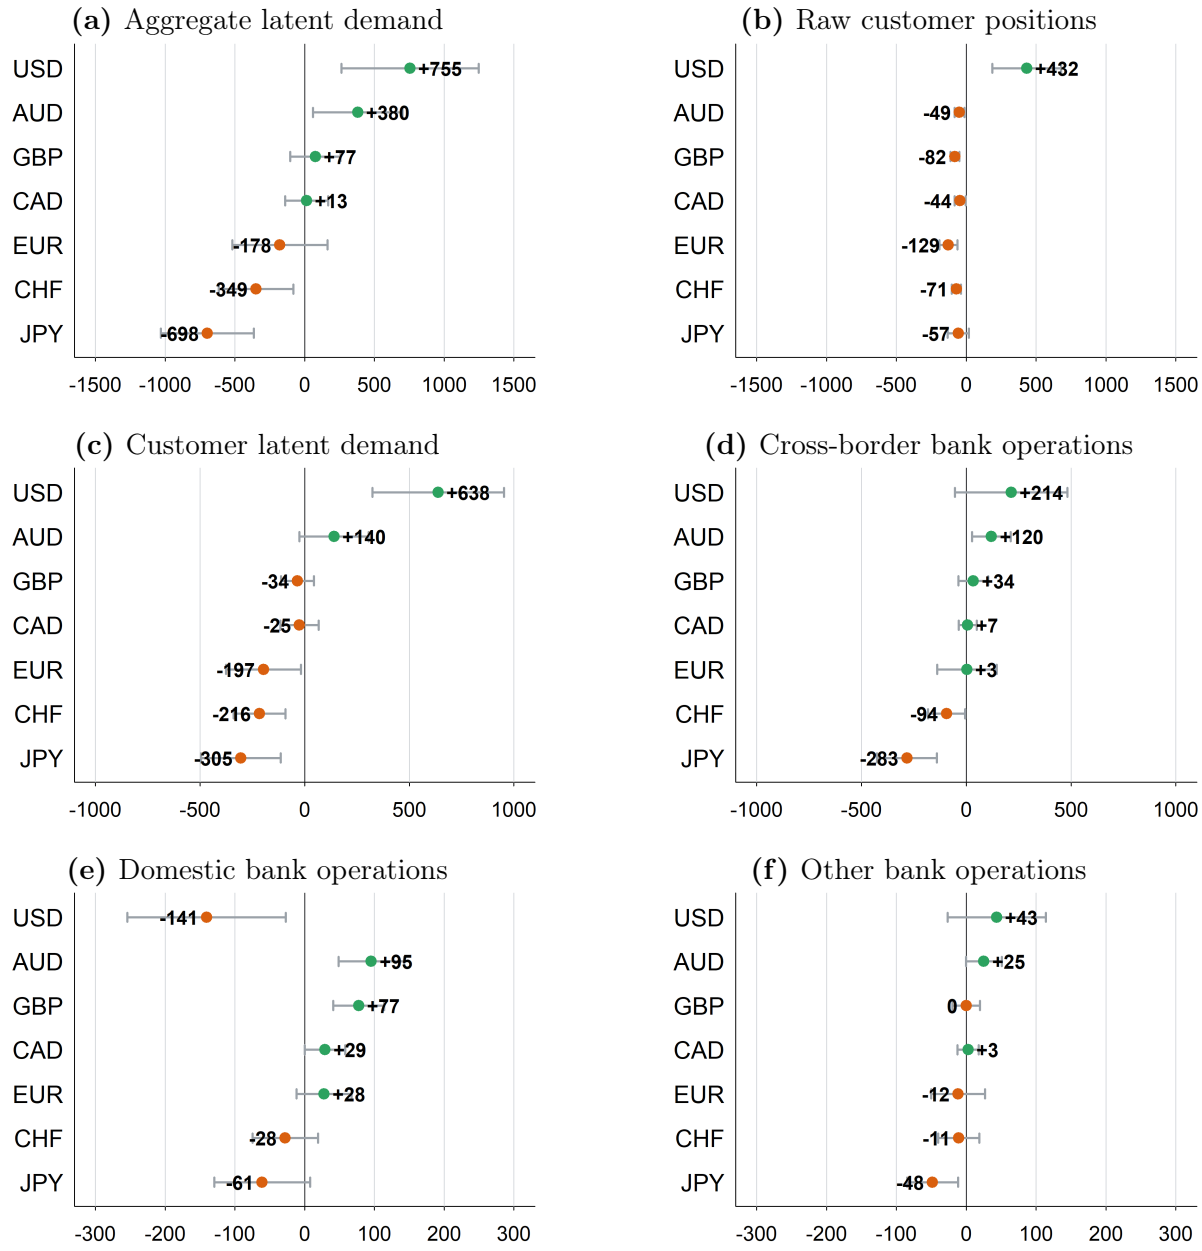
The composition of $\mathbf{A}_t$ differs substantially from the distribution of investor size. Customers and cross-border bank operations account for 38.9 and 38.0 percent of total aggregate position responsiveness, respectively. The contribution of cross-border bank operations is particularly notable because they account for only 17.6 percent of investor size, reflecting their much higher estimated price responsiveness.

Figure 5a shows positive aggregate latent demand for USD and AUD and negative aggregate latent demand for JPY, CHF, and EUR. Figure 5b provides a useful comparison: raw customer positions are much more dollar-centric, with customers demanding USD and supplying every other currency on average. Aggregate latent demand is therefore not simply a relabeling of observed customer positions. It combines latent demand across customers and bank operations after removing the price-induced component of their positions.

Figures 5c, 5d, 5e, and 5f decompose aggregate latent demand by investor group. Customers generate most of the average USD demand. Within the banking sector, cross-border bank operations are the largest source of USD demand and JPY supply, whereas domestic bank operations supply USD on average and offset part of the cross-border demand.

The two equilibrium objects enter the counterfactuals separately: removing an investor changes both the latent demand it contributes and the aggregate position responsiveness available to absorb the remaining demand.

Figure 5: Aggregate Latent Demand by Currency and Investor Group



Notes: The aggregate latent-demand and investor-group subfigures report the time-series mean of latent demand by currency, with horizontal intervals extending one monthly standard deviation on either side of the mean. The raw-customer-position subfigure reports the corresponding mean and standard deviation of aggregate observed customer positions before separating the price-induced component. Positive values indicate net demand and negative values indicate net supply. Aggregate latent demand equals the sum of the four investor-group components. Values are in USD billions. Subfigures in the same row use the same horizontal scale; scales differ across rows.

6 Counterfactuals

The estimated demand system allows us to study how the FX funding market responds to disruptions in financial intermediation. We consider two distinct counterfactuals with direct

real-world counterparts. Section 6.1 removes individual bank operations from the market entirely, as may occur following an operational outage, bank distress, resolution, or broad balance-sheet retrenchment. Section 6.2 instead considers a U.S. policy intervention that prohibits bank operations connected to the United States through headquarters or physical location from supplying a target currency while allowing them to remain active in other currencies.

Both exercises require an equilibrium model because the impact of a disruption depends not only on the positions directly affected, but also on how customers and the remaining bank operations respond as funding prices change. The relevant measure of vulnerability is therefore not the normal-time size of an affected operation or sanction perimeter alone, but the availability of other investors willing to replace the lost intermediation. The two exercises implement this idea differently. The withdrawal counterfactual removes both an operation's entire position and its contribution to the aggregate position-response matrix; for any fixed withdrawal set, the resulting equilibrium is linear. The sanctions counterfactual leaves covered operations active and instead solves a nonlinear constrained equilibrium because whether a bank is prevented from supplying the target currency depends on counterfactual prices.

6.1 Resilience to the Loss of Dollar Funding Intermediaries

We first ask how resilient Dollar funding is to the loss of individual bank operations. Such disruptions are typically institution-specific. A bank may temporarily lose access to the market because of an operational outage or cyber incident, experience funding stress that forces it to retrench, or become unable to intermediate while undergoing resolution. The immediate question for the market is whether other intermediaries can step in.

We identify bank operations that are net suppliers of USD on average and rank them from largest to smallest according to their average Dollar supply. Let $\mathcal{G}_k$ denote the first k operations in this ranking. At point k of the counterfactual path, all operations in $\mathcal{G}_k$ exit the market entirely: their full seven-currency positions and their contributions to aggregate price responsiveness are removed.

For any cumulative withdrawal set $\mathcal{G}$, let $\mathbf{A}_{\mathcal{G},t}$ denote the operations' contribution to the aggregate position-response matrix and let $\mathbf{Q}_{\mathcal{G},t}$ denote their aggregate factual currency-position vector. The counterfactual price change satisfies

$$(\mathbf{A}_t - \mathbf{A}_{\mathcal{G},t})\left(\mathbf{P}_t^{\mathcal{G}} - \mathbf{P}_t\right) = -\mathbf{Q}_{\mathcal{G},t}. \tag{16}$$

The equation has a simple interpretation. Removing $\mathcal{G}$ leaves the surviving investors to

replace the withdrawn position vector $\mathbf{Q}_{\mathcal{G},t}$, while $\mathbf{A}_t - \mathbf{A}_{\mathcal{G},t}$ measures the price responsiveness that remains available to do so. Price effects can therefore grow for two reasons: a larger position must be replaced, and the market has less remaining capacity to absorb it. Appendix B.1 derives equation (16) from the factual and counterfactual market-clearing conditions.

Figure 6 plots the cumulative effect on the USD funding benchmark. The solid line reports the equilibrium effect obtained from equation (16) for each cumulative withdrawal set. The dashed line instead adds the one-operation effects obtained by removing each operation separately from the intact market. If the importance of each operation were independent of which other intermediaries remained available, the two curves would coincide. Their difference, shown in the lower portion of the figure, measures the nonlinearity created by the loss of replacement capacity.

The first removals illustrate the market's resilience. Several of the largest normal-time Dollar suppliers are domestic bank operations. Their disappearance raises the funding price, but the effect remains relatively contained because other banks and customers adjust their positions in response. The system therefore does not rely mechanically on the continued presence of its largest suppliers.

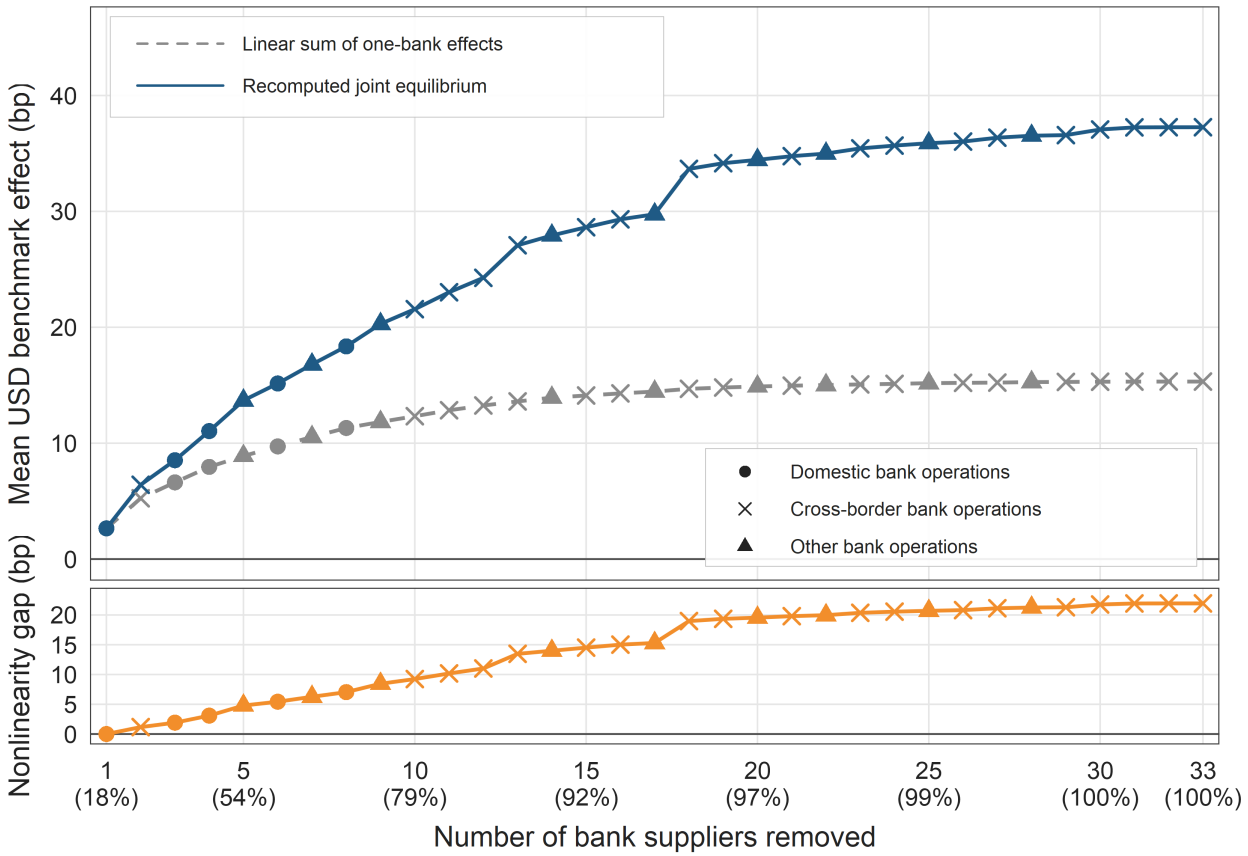
The larger increases later in the path come as cross-border bank operations are removed. These operations are not necessarily the largest suppliers in the intact market, but their relatively high price responsiveness makes them important sources of replacement capacity. Once earlier suppliers have disappeared, removing these potential backstops sharply reduces the responsiveness of the surviving market and produces larger price increases.

This distinction is central to financial resilience. The institutions that make a market resilient need not be those with the largest normal-time market shares. Resilience depends on the presence of intermediaries that can expand when others are impaired. Cross-border bank operations play precisely this role in Dollar funding: their importance comes not only from the positions they normally hold, but also from their ability to accommodate funding pressure when the system is under stress.

The widening gap between the solid and dashed lines captures this deterioration in resilience. After all 33 bank operations that supply USD on average have been removed, the recomputed joint-equilibrium effect reaches 37.3 basis points, compared with 15.3 basis points from summing their separate effects in the intact market. The linear sum understates the joint effect because each one-operation calculation assumes that the rest of the banking network remains available, whereas the cumulative counterfactual progressively removes the operations that would otherwise replace one another.

Because each operation exits completely, equation (16) depends on its entire seven-

Figure 6: Resilience to the Loss of Dollar Funding Intermediaries



Notes: Bank operations that are net suppliers of USD on average are ranked from largest to smallest average supply and removed cumulatively. The solid line reports the recomputed joint-equilibrium effect, while the dashed line sums one-operation effects measured in the intact market. The lower portion reports the difference between the two. The symbol at each step identifies the operation removed: circles denote domestic bank operations, crosses denote cross-border bank operations, and triangles denote other bank operations. Percentages below selected horizontal-axis values are cumulative shares of average bank USD supply removed. The USD funding benchmark is measured relative to the equal-weight average of the other six currencies.

currency position rather than only its USD supply. Other currency legs may reinforce or offset the direct Dollar effect, so individual steps need not be monotonic. The broader message is that normal-time market share alone is not a sufficient measure of resilience. Dollar funding can absorb isolated disruptions while becoming much more fragile when disruptions reach several of the intermediaries that would otherwise serve as substitutes.

6.2 Unilateral and Coordinated Financial Sanctions

The previous exercise considers disruptions to individual intermediaries. Governments, however, can also restrict financial intermediation directly. Financial sanctions may prohibit

banks subject to a jurisdiction's authority from providing particular forms of funding or transacting in particular currencies. We use the model to ask how much influence the United States can exert over a currency's equilibrium funding conditions through such a restriction.

For each target currency n, we consider two U.S. sanction scopes. The "HQ only" sanction covers all bank operations headquartered in the United States. The "HQ plus location" sanction covers the union of bank operations headquartered in the United States and bank operations physically located there, counting operations that satisfy both conditions only once. The broader scope therefore encompasses both banking organizations controlled from the United States and foreign banking organizations operating within it.

The sanction applies only to the target currency. Every covered bank operation is prohibited from being a net supplier of currency n in the counterfactual equilibrium, but remains active and price-responsive in all other currencies. When the prohibition binds, the operation's target-currency position is set to zero and its positive non-target positions are reduced proportionally to preserve adding up. Because the set of binding operations depends on counterfactual prices, the restrictions and relative prices are determined simultaneously. The experiment therefore represents a restriction on one intermediation activity rather than the disappearance of the bank. Appendix B.2 provides the formal construction.

We compare each U.S. sanction with a coordinated benchmark in which the same hard no-supply restriction applies to all 48 bank operations, including domestic, cross-border, and other bank operations. Customers remain unconstrained. Let ΔP_n^s denote the change in the target-currency funding benchmark under U.S. sanction scope $s \in \{\text{HQ}, \text{HQ} + \text{Loc}\}$, and let $\Delta P_n^{\text{Coord}}$ denote the corresponding change under the coordinated restriction. We summarize U.S. financial leverage by

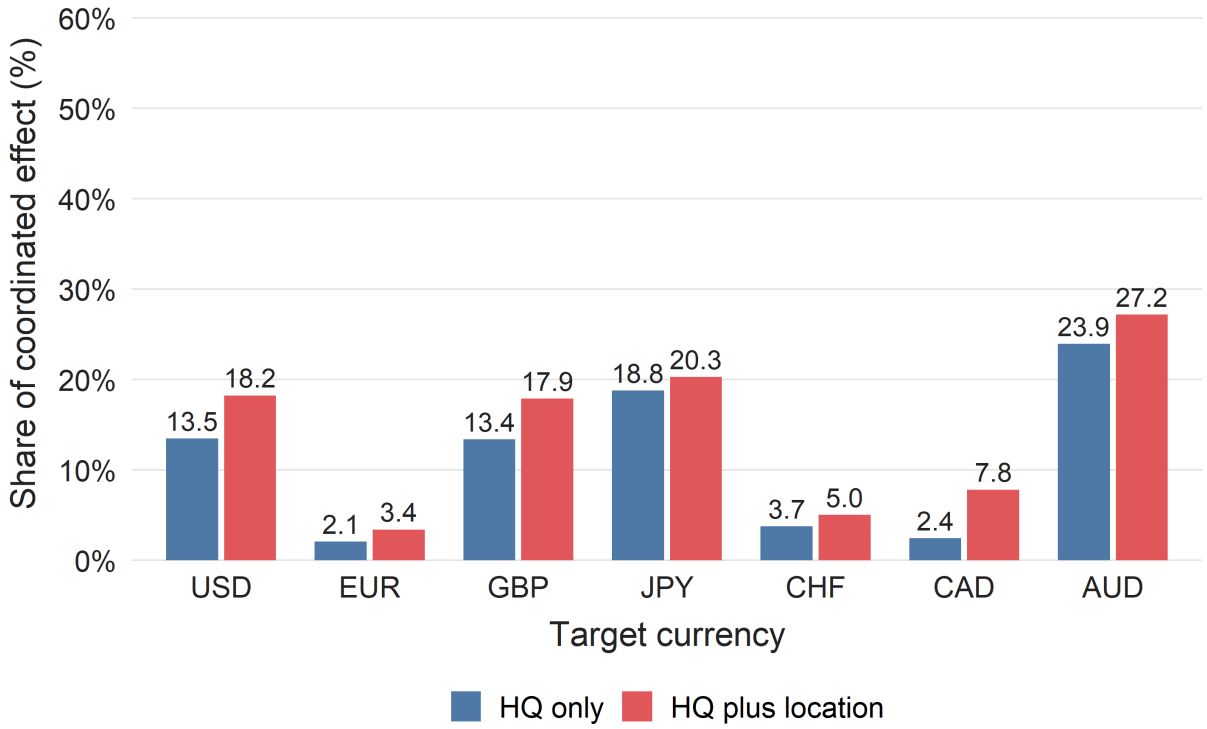
$$\mathcal{L}_n^s = \frac{\Delta P_n^s}{\Delta P_n^{\text{Coord}}}, \qquad s \in \{\text{HQ}, \text{HQ} + \text{Loc}\}. \tag{17}$$

The ratio measures how much of the funding disruption attainable under coordinated sanctions can be generated by the United States acting alone.

Figure 7 reports these ratios across the seven target currencies. For each currency, the coordinated benchmark is normalized to 100 percent. The two bars distinguish the reach of the United States through bank operations headquartered there from its broader reach through bank operations headquartered in or physically located there.

The comparison separates the international presence of the U.S. banking system from its ability to move market prices. The United States may control large banking organizations and host substantial FX activity without possessing a corresponding choke point over every currency's funding market. If bank operations outside the U.S. sanction perimeter can

Figure 7: U.S. Financial Leverage Relative to Coordinated Sanctions



Notes: For each target currency, the bars report the U.S. unilateral equilibrium effect as a percentage of the coordinated effect for the same currency. "HQ only" covers bank operations headquartered in the United States. "HQ plus location" covers the union of bank operations headquartered in and physically located in the United States. The coordinated benchmark applies the hard no-supply restriction to all 48 bank operations, including domestic, cross-border, and other bank operations; customers remain unconstrained.

replace the lost supply, the equilibrium price effect remains small even when the covered banks initially account for substantial activity.

Physical location expands the reach of the U.S. sanction. Restricting U.S.-headquartered bank operations alone generates 13.5 percent of the coordinated USD disruption. Expanding the sanction to bank operations headquartered in or physically located in the United States raises this share to 18.2 percent. The broader U.S. perimeter generates 27.2 percent of the coordinated AUD disruption, 20.3 percent of the JPY disruption, and 17.9 percent of the GBP disruption, while its leverage over EUR, CHF, and CAD remains considerably smaller.

The USD result is economically meaningful but well below the coordinated benchmark. Even under the broader "HQ plus location" scope, the United States acting alone reproduces less than one fifth of the disruption generated when all bank operations are prohibited from supplying USD. Dollar funding is therefore exposed to a unilateral U.S. restriction, but its supply is distributed across a multinational banking network rather than controlled

exclusively by U.S.-headquartered or U.S.-located banks.

More generally, a large value of $\mathcal{L}_n^s$ identifies a currency whose intermediation network lies disproportionately within the reach of the United States. A small value indicates that unaffected foreign banks can replace much of the restricted supply. The comparison across currencies therefore reveals an important dimension of financial resilience: whether currency funding is internationally diversified across banking systems.

Taken together, the two counterfactuals show that resilience in global funding markets depends critically on the availability of alternative intermediaries. Following an institution-specific disruption, surviving banks can expand their positions and become marginal backstops. The same mechanism limits the effectiveness of U.S. financial sanctions when bank operations outside the U.S. sanction perimeter remain free to intermediate. Large disruptions arise only when restrictions reach sufficiently broadly across the institutions that would otherwise provide this substitution.

7 Conclusion

This paper studies the resilience of global financial intermediation through the lens of the FX funding market. Using granular data that distinguish customers from banks and, for banks, headquarters from physical location, we document a market in which customers operate in strong home-currency habitats while banks intermediate substantial funding through cross-border operations. We estimate an equilibrium model that allows investors to differ in their currency anchors and responsiveness to funding prices. Estimation reveals that cross-border bank operations are substantially more price responsive than other investors, making them particularly important adjustment margins when the market is disrupted. Our counterfactuals show that this substitution makes FX funding resilient to isolated disruptions. Surviving banks expand to replace lost intermediation, and cross-border operations can become important backstops even when their normal-time positions are modest. The same mechanism limits unilateral financial leverage. Restrictions imposed by a single jurisdiction generate only a fraction of the price effects achieved through coordinated sanctions because banks elsewhere step in.

More broadly, our results illustrate why the vulnerability of a financial system cannot be inferred from the size or normal-time activity of its participants alone. What matters is not only what an institution does today, but who can perform that function if the institution becomes unavailable, and whether those substitutes remain available when disruptions become broader. International financial intermediation provides considerable resilience precisely because activities can migrate across institutions and jurisdictions. Yet this resilience

is not immutable. As disruptions engulf the intermediaries that would otherwise step in, the system can become increasingly difficult to stabilize. Understanding financial resilience therefore requires mapping not only where intermediation resides, but where its substitutes reside. In an increasingly fragmented world, the robustness of global finance may ultimately depend on preserving those margins of substitution.

References

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi. 2015. Systemic risk and stability in financial networks. *American Economic Review* 105:564–608.

Allen, F., and D. Gale. 2000. Financial contagion. *Journal of Political Economy* 108:1–33.

Cabrales, A., P. Gottardi, and F. Vega-Redondo. 2017. Risk-sharing and contagion in networks. *Review of Financial Studies* 30:3086–127.

Cetorelli, N., and L. S. Goldberg. 2012. Banking globalization and monetary transmission. *Journal of Finance* 67:1811–43.

Du, W., B. Hébert, and A. W. Huber. 2023. Are intermediary constraints priced? *Review of Financial Studies* 36:1464–507.

Du, W., and A. Huber. 2024. Dollar asset holdings and hedging around the globe. Working Paper.

Du, W., A. Tepper, and A. Verdelhan. 2018. Deviations from covered interest rate parity. *Journal of Finance* 73:915–57.

Elliott, M., B. Golub, and M. O. Jackson. 2014. Financial networks and contagion. *American Economic Review* 104:3115–53.

Gabaix, X., and R. S. Koijen. 2024. Granular instrumental variables. *Journal of Political Economy* 132:2274–303.

Huber, A., and S. Kundu. 2026. When funding markets move credit markets: Foreign investors and u.s. clos. Working Paper.

Khetan, U. 2026. Synthetic dollar funding. Working Paper.

Kloks, P., E. Mattille, and A. Ranaldo. 2026. Segmented dollar funding. *Journal of Financial Economics* 184:104348–.

Kubitza, C., J.-D. Sigaux, and Q. Vandeweyer. 2024. The implications of cip deviations for international capital flows. *Fama-Miller Working Paper, Chicago Booth Research Paper* .

Nenova, T., A. Schrimpf, and H. S. Shin. 2025. Global portfolio investments and fx derivatives. BIS Working Papers 1273, Bank for International Settlements.

Schnabl, P. 2012. The international transmission of bank liquidity shocks: Evidence from an emerging market. *Journal of Finance* 67:897–932.

A Estimation Details

This appendix provides the implementation details underlying Section 5. Throughout, t indexes calendar months, while τ indexes observed dates within a month.

A.1 Monthly Panel and Target Blocks

For each calendar month t, let $\mathcal{D}_t$ denote the first five dates on which the complete seven-currency price vector is finite. Monthly positions and prices are averaged over exactly the same dates:

$$Q_{i,n,t} := \frac{1}{5}\sum_{\tau\in\mathcal{D}_t} Q_{i,n,\tau}, \qquad P_{n,t} := \frac{1}{5}\sum_{\tau\in\mathcal{D}_t} P_{n,\tau}. \tag{A1}$$

If an earlier date has an incomplete price vector, it is skipped for both positions and prices. We then construct the size-normalized position and its monthly change:[7]

$$S_{i,t} := \sum_{n=0}^{N} |Q_{i,n,t}|, \qquad q_{i,n,t} := \frac{Q_{i,n,t}}{S_{i,t-1}}, \qquad \Delta q_{i,n,t} := q_{i,n,t} - q_{i,n,t-1}. \tag{A2}$$

The currency-anchor assignments described in Section 3 are computed from full-sample position levels and then held fixed. The estimation sample contains eight active customer regions and 48 active bank operations. The bank operations comprise 7 domestic operations, 24 cross-border operations, and 17 other operations, accounting for 43.6, 31.1, and 25.3 percent of average bank size, respectively.

Each customer region contributes six non-anchor currency equations. For bank operations, raw positions and size are aggregated within operation group g and assigned currency anchor h before normalization. Let $\mathcal{I}_{g,h}$ denote the bank operations in group $g \in \{D, X, O\}$ assigned to currency anchor h. We construct

$$Q^B_{g,h,n,t} := \sum_{i\in\mathcal{I}_{g,h}} Q_{i,n,t}, \qquad S^B_{g,h,t} := \sum_{i\in\mathcal{I}_{g,h}} S_{i,t}, \qquad q^B_{g,h,n,t} := \frac{Q^B_{g,h,n,t}}{S^B_{g,h,t-1}}. \tag{A3}$$

This preserves the bilateral price associated with each assigned currency anchor and avoids equally averaging already normalized bank operations. Customer region–currency blocks receive equal weight; within each bank-operation group, currency-anchor-by-non-anchor-currency blocks also receive equal weight.

[7] The size measure $S_{i,t}$ is twice the one-sided gross position used in Section 3; this constant scaling only fixes the units of normalized positions and price-responsiveness estimates.

A.2 Construction of the Granular Instruments

We construct currency-level granular instruments separately from bank and customer position changes. Let $j \in \{B, C\}$ denote the source side and let $r = 1, \ldots, J_j$ index source units. Source units are ranked from largest to smallest by average lagged size. We retain the largest units until their cumulative share reaches 99 percent of source-side activity and combine any remaining units into one residual unit. This produces 38 bank source units—37 bank operations and one residual unit—and eight customer source units.

For each source side and currency, we remove the equal-weight cross-sectional mean from normalized position changes:

$$\Delta\overline{q}^{j}_{n,t} := \frac{1}{J_j}\sum_{r=1}^{J_j}\Delta q^{j}_{r,n,t}, \qquad \Delta\check{q}^{j}_{r,n,t} := \Delta q^{j}_{r,n,t} - \Delta\overline{q}^{j}_{n,t}. \tag{A4}$$

After replacing missing source-unit changes by their time-series means, we standardize each source-unit series and estimate a separate correlation-PCA factor model for each currency:

$$\Delta\check{q}^{j}_{r,n,t} = \boldsymbol{\lambda}^{j\top}_{r,n}\mathbf{f}^{j}_{n,t} + u^{j}_{r,n,t}. \tag{A5}$$

We retain six factors for bank-origin instruments and one factor for customer-origin instruments, using the same number of factors across currencies within each source side. The estimated residuals are transformed back to the original units before aggregation.

Let the lagged size share of source unit r be

$$\omega^{j}_{r,t-1} := \frac{S^{j}_{r,t-1}}{\sum_{\ell=1}^{J_j} S^{j}_{\ell,t-1}}. \tag{A6}$$

The currency-level granular instrument is the size-weighted sum of estimated idiosyncratic shocks:

$$G^{j}_{n,t} := \sum_{r=1}^{J_j}\omega^{j}_{r,t-1}\widehat{u}^{j}_{r,n,t}. \tag{A7}$$

Thus, equal weights are used to remove the cross-sectional mean, while lagged size weights convert source-unit shocks into an aggregate currency-level shift. For a bilateral equation comparing currency n with currency anchor h, the primitive instrument is

$$z^{j}_{h,n,t} := G^{j}_{n,t} - G^{j}_{h,t}. \tag{A8}$$

Customer equations use bank-origin instruments, while bank-operation equations use customer-

origin instruments. The source-side PCA factors for currencies n and h enter separately as controls in both stages.

A.3 Cross-Fitted IV Estimation and Inference

Let

$$x_{h,n,t} := \Delta(P_{n,t} - P_{h,t}) \tag{A9}$$

denote the endogenous bilateral price change, and let $\mathbf{c}^j_{h,n,t}$ collect an intercept and the source-side factors for currencies n and h.

To combine the bilateral instruments while avoiding in-sample estimation of their price pass-through, we divide the monthly estimating sample into five contiguous folds. For each currency pair and held-out fold b, the other four folds are used to estimate

$$x_{h,n,t} = \pi^{(-b)}_{h,n} z^j_{h,n,t} + \mathbf{c}^{j\top}_{h,n,t} \mathbf{a}^{(-b)}_{h,n} + v_{h,n,t}, \qquad t \notin \mathcal{T}_b. \tag{A10}$$

Using the same training observations, we regress the primitive instrument on the factor controls and denote the resulting control coefficient by $\widehat{\mathbf{b}}^{(-b)}_{h,n}$. For an observation in the held-out fold, the scalar instrument is

$$z^{*,j}_{h,n,t} := \widehat{\pi}^{(-b)}_{h,n} \left(z^j_{h,n,t} - \mathbf{c}^{j\top}_{h,n,t} \widehat{\mathbf{b}}^{(-b)}_{h,n} \right), \qquad t \in \mathcal{T}_b. \tag{A11}$$

Repeating this procedure across folds gives a scalar instrument for every pair-month observation, with the pair-specific pass-through weight and control projection estimated outside that observation's fold. The PCA factors and primitive granular instruments are constructed using the full monthly sample; cross-fitting applies to the pair-specific pass-through weights and control projections.

Let a index a target portfolio within group g: an individual customer region when $g = C$, and a currency-anchor-specific bank portfolio when $g \in \{D, X, O\}$. Let $h(a)$ denote the portfolio's currency anchor and $y^g_{a,n,t}$ its normalized position change. For each target group g, we pool the relevant bilateral blocks and estimate

$$y^g_{a,n,t} = -\theta_g x_{h(a),n,t} + \mathbf{c}^{j\top}_{h(a),n,t} \boldsymbol{\gamma}^g_{h(a),n} + \varepsilon^g_{a,n,t}, \qquad n \neq h(a), \tag{A12}$$

using $z^{*,j}_{h(a),n,t}$ as an instrument for $x_{h(a),n,t}$. Price responsiveness is common across bilateral blocks within group g, while the control coefficients vary by currency pair. Inference is conducted at the monthly level, allowing observations from different bilateral blocks in the same month to be correlated. We use Newey–West standard errors with lag length selected

by the rule of thumb $L_T = \lfloor 4(T/100)^{2/9} \rfloor$, where T is the number of monthly observations. This rule selects four lags in our sample. The reported first-stage F-statistics use the same HAC adjustment.

A.4 Derivation of the Aggregate Equilibrium

Let $\mathbf{q}_{i,t}^{\mathrm{NH}}$ and $\widehat{\mathbf{d}}_{i,t}$ collect investor i's normalized positions and recovered latent demands for its N non-anchor currencies. Using the anchor-incidence matrix defined in Section 5.3, the investor-level demand system can be written as

$$\mathbf{q}_{i,t}^{\mathrm{NH}} = \widehat{\mathbf{d}}_{i,t} - \widehat{\theta}_{g(i)} \mathbf{H}_{h(i)} \mathbf{P}_t. \tag{A13}$$

For any vector of non-anchor positions $\mathbf{x}$, the transformation $\mathbf{H}_{h(i)}^{\top}\mathbf{x}$ places each component in its corresponding non-anchor currency and assigns the negative of their sum to the currency anchor. It therefore produces a full currency-position vector whose elements add up to zero. Investor i's position vector in position units is consequently

$$\mathbf{Q}_{i,t} = S_{i,t-1} \mathbf{H}_{h(i)}^{\top} \widehat{\mathbf{d}}_{i,t} - S_{i,t-1} \widehat{\theta}_{g(i)} \mathbf{H}_{h(i)}^{\top} \mathbf{H}_{h(i)} \mathbf{P}_t. \tag{A14}$$

Summing across investors gives

$$\sum_i \mathbf{Q}_{i,t} = \sum_i S_{i,t-1} \mathbf{H}_{h(i)}^{\top} \widehat{\mathbf{d}}_{i,t} - \left(\sum_i S_{i,t-1} \widehat{\theta}_{g(i)} \mathbf{H}_{h(i)}^{\top} \mathbf{H}_{h(i)} \right) \mathbf{P}_t = \mathbf{b}_t - \mathbf{A}_t \mathbf{P}_t. \tag{A15}$$

Market clearing sets $\sum_i \mathbf{Q}_{i,t} = 0$, yielding

$$\mathbf{A}_t \mathbf{P}_t = \mathbf{b}_t, \tag{A16}$$

where $\mathbf{b}_t$ and $\mathbf{A}_t$ are defined in equation (15). Because only relative funding prices matter, we impose $P_{\mathrm{USD},t} = 0$ when solving the system.

B Counterfactual Computation

This appendix provides the computational details for the two counterfactual exercises in Section 6. The first derives the equilibrium response when selected USD funding suppliers exit the market entirely, while the second describes the constrained equilibrium in which bank operations within the U.S. sanction perimeter remain active but are prohibited from supplying a target currency.

B.1 Cumulative Withdrawal of USD Funding Suppliers

Let $\overline{Q}_{i,\mathrm{USD}} := T^{-1}\sum_t Q_{i,\mathrm{USD},t}$, and order the bank operations satisfying $\overline{Q}_{i,\mathrm{USD}} < 0$ from largest to smallest average USD supply. Let $\mathcal{G}_k$ denote the first k operations in this ranking. Customers are never included in the withdrawal set, and each selected bank operation exits its entire seven-currency funding book.

For an arbitrary withdrawal set $\mathcal{G}$, define its contributions to aggregate latent demand and the aggregate position-response matrix as

$$\mathbf{b}_{\mathcal{G},t} := \sum_{i\in\mathcal{G}} \mathbf{b}_{i,t}, \qquad \mathbf{A}_{\mathcal{G},t} := \sum_{i\in\mathcal{G}} \mathbf{A}_{i,t}. \tag{A17}$$

Its factual position vector is

$$\mathbf{Q}_{\mathcal{G},t} = \mathbf{b}_{\mathcal{G},t} - \mathbf{A}_{\mathcal{G},t}\mathbf{P}_t. \tag{A18}$$

Before the withdrawal, equilibrium satisfies $\mathbf{A}_t\mathbf{P}_t = \mathbf{b}_t$. After the operations in $\mathcal{G}$ exit, both their latent demand and their contribution to price responsiveness disappear, so

$$(\mathbf{A}_t - \mathbf{A}_{\mathcal{G},t})\,\mathbf{P}_t^{\mathcal{G}} = \mathbf{b}_t - \mathbf{b}_{\mathcal{G},t}. \tag{A19}$$

Subtracting the factual equilibrium, expressed using the surviving position-response matrix, gives

$$(\mathbf{A}_t - \mathbf{A}_{\mathcal{G},t})\left(\mathbf{P}_t^{\mathcal{G}} - \mathbf{P}_t\right) = (\mathbf{b}_t - \mathbf{b}_{\mathcal{G},t}) - (\mathbf{A}_t - \mathbf{A}_{\mathcal{G},t})\,\mathbf{P}_t = -\mathbf{Q}_{\mathcal{G},t}, \tag{A20}$$

which is equation (16). The right-hand side is the full seven-currency position that disappears, while the matrix on the left-hand side measures the price responsiveness remaining among surviving investors. As the withdrawal set grows, more positions must be replaced and less responsiveness remains available to absorb them.

For any price-change vector $\Delta\mathbf{P}_t$, define the change in the USD funding benchmark as

$$\mathcal{B}_{\mathrm{USD}}\left(\Delta\mathbf{P}_t\right) := \Delta P_{\mathrm{USD},t} - \frac{1}{6}\sum_{n\neq\mathrm{USD}} \Delta P_{n,t}. \tag{A21}$$

The solid line in Figure 6 reports the time-series average of $\mathcal{B}_{\mathrm{USD}}(\mathbf{P}_t^{\mathcal{G}_k} - \mathbf{P}_t)$. For the dashed line, each operation is instead removed separately from the intact market by solving

$$(\mathbf{A}_t - \mathbf{A}_{i,t})\,\Delta\mathbf{P}_{i,t}^{\mathrm{one}} = -\mathbf{Q}_{i,t}, \tag{A22}$$

and the resulting average one-operation effects are summed through rank k. The lower portion of the figure reports the joint effect minus this linear sum.

B.2 U.S. Hard-Prohibition Financial Sanctions

Let $\mathcal{B}$ denote the set of bank operations, and let $\mathrm{HQ}(i)$ and $\mathrm{Loc}(i)$ denote operation i's headquarters and physical location. The two U.S. coverage sets are

$$\mathcal{C}^{\mathrm{HQ}} := \{i \in \mathcal{B} : \mathrm{HQ}(i) = \mathrm{US}\}, \qquad \mathcal{C}^{\mathrm{HQ+Loc}} := \{i \in \mathcal{B} : \mathrm{HQ}(i) = \mathrm{US} \text{ or } \mathrm{Loc}(i) = \mathrm{US}\}. \tag{A23}$$

Thus, HQ plus location is the union of the two sets rather than the sum of two separate sanctions; an operation satisfying both conditions is included once. The coordinated coverage set is $\mathcal{C}^{\mathrm{Coord}} = \mathcal{B}$, which contains all 48 bank operations. Customers are outside every coverage set.

Let $\mathbf{b}_{i,t}$ and $\mathbf{A}_{i,t}$ denote investor i's contributions to aggregate latent demand and the aggregate position-response matrix. At a candidate price vector $\mathbf{P}$, its unconstrained full currency-position vector is

$$\mathbf{Q}^u_{i,t}(\mathbf{P}) = \mathbf{b}_{i,t} - \mathbf{A}_{i,t}\mathbf{P}. \tag{A24}$$

Fix a target currency n and a coverage set $\mathcal{C}$. If $i \notin \mathcal{C}$, or if $Q^u_{i,n,t}(\mathbf{P}) \geq 0$, the operation keeps its unconstrained position. If $i \in \mathcal{C}$ and $Q^u_{i,n,t}(\mathbf{P}) < 0$, the hard prohibition binds. Define the positive non-target counter-leg set

$$\mathcal{D}_{i,n,t}(\mathbf{P}) := \{k \neq n : Q^u_{i,k,t}(\mathbf{P}) > 0\} \tag{A25}$$

and the proportional contraction factor

$$\lambda_{i,n,t}(\mathbf{P}) := 1 - \frac{-Q^u_{i,n,t}(\mathbf{P})}{\sum_{k \in \mathcal{D}_{i,n,t}(\mathbf{P})} Q^u_{i,k,t}(\mathbf{P})}. \tag{A26}$$

The constrained position is

$$\widetilde{Q}^{\mathcal{C},n}_{i,k,t}(\mathbf{P}) = \begin{cases} 0, & k = n, \\ \lambda_{i,n,t}(\mathbf{P}) Q^u_{i,k,t}(\mathbf{P}), & k \in \mathcal{D}_{i,n,t}(\mathbf{P}), \\ Q^u_{i,k,t}(\mathbf{P}), & \text{otherwise.} \end{cases} \tag{A27}$$

The contraction removes an amount of positive non-target positions equal to the prohibited target-currency supply, so the constrained position continues to add up to zero. Other negative non-target positions are unchanged. The rule is applied to every covered operation simultaneously.

For each month, target currency, and coverage set, we solve for the six independent

relative prices such that the change in aggregate positions is zero:

$$\sum_i \left[\widetilde{\mathbf{Q}}_{i,t}^{\mathcal{C},n}(\mathbf{P}_t^{\mathcal{C},n}) - \mathbf{Q}_{i,t}^{u}(\mathbf{P}_t)\right] = \mathbf{0}, \qquad P_{\mathrm{USD},t}^{\mathcal{C},n} = P_{\mathrm{USD},t} = 0. \tag{A28}$$

For uncovered investors, $\widetilde{\mathbf{Q}}_{i,t}^{\mathcal{C},n} = \mathbf{Q}_{i,t}^{u}$. Because both the unconstrained positions and the set of binding covered operations depend on the candidate prices, equation (A28) is nonlinear. At every trial price vector, the code recomputes all investors' unconstrained positions, reapplies the prohibition to all covered operations, and updates the six relative prices until equation (A28) holds. Thus, a covered bank that is initially a demander but would become a supplier as prices move is also constrained in equilibrium, while a covered bank that remains a demander is unrestricted.

For any price vector $\mathbf{P}$, define the funding benchmark for target currency n as

$$\mathcal{P}_n(\mathbf{P}) := P_n - \frac{1}{N}\sum_{k\neq n} P_k. \tag{A29}$$

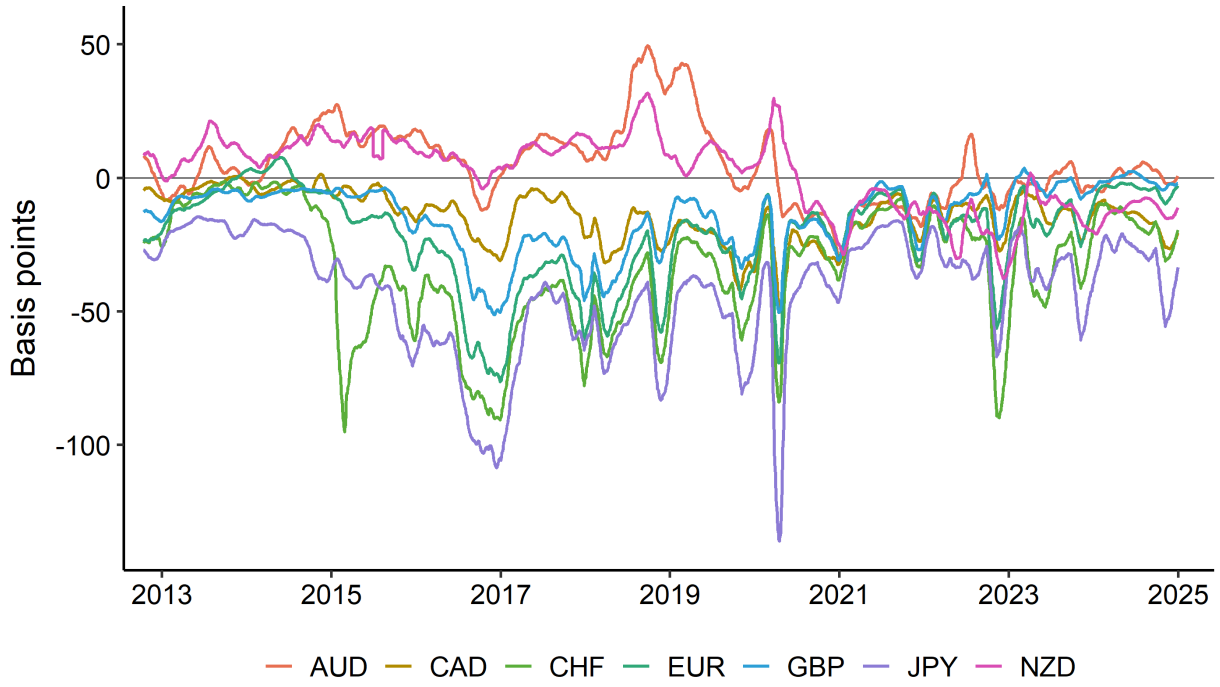
For U.S. sanction scope $s \in \{\mathrm{HQ}, \mathrm{HQ}+\mathrm{Loc}\}$, the reported unilateral effect is

$$\Delta P_n^s := \frac{1}{T}\sum_t \left[\mathcal{P}_n\left(\mathbf{P}_t^{\mathcal{C}^s,n}\right) - \mathcal{P}_n(\mathbf{P}_t)\right]. \tag{A30}$$

The coordinated effect $\Delta P_n^{\mathrm{Coord}}$ is defined analogously using $\mathcal{C}^{\mathrm{Coord}} = \mathcal{B}$. Figure 7 reports $100\mathcal{L}_n^s$, where $\mathcal{L}_n^s$ is defined in equation (17).

C Additional Figures and Tables

Figure A1: Price of FX funding across currencies at 3-month tenor



Notes: This figure plots $P_{n,t,3\text{M}}$, the 3M price of obtaining funding in currency n by supplying USD funding for each non-USD currency n. Positive values indicate that borrowing USD synthetically via currency n is cheaper than direct USD funding; negative values indicate the reverse. The series show 30-day moving averages.