Funding Gravity and Financial Power*

Yu An† Amy W. Huber‡

This version: October 2025.
Most recent version here.

Abstract

We quantify countries' financial power using funding gravity: a measure of a country's ability to attract capital when its own financial conditions change. Financial power shapes global capital flows, yet it has proven difficult to measure directly. We construct a novel empirical measure of funding gravity in the global foreign exchange (FX) funding market using proprietary data and a novel methodology that links observed funding prices, measured by deviations from covered interest parity (CIP), to the corresponding funding quantities. To address the endogeneity created by correlated shocks across countries, we identify funding blocs, where shocks within each bloc affect member countries according to their participation intensity, while shocks across different blocs are distinct and orthogonal. A country's funding gravity is therefore determined by its participation across blocs and by each bloc's intrinsic ability to attract capital. Our estimates uncover a striking hierarchy of financial power. The United States' funding gravity equals that of all other countries combined. The euro area commands roughly half the U.S. level, while the United Kingdom and Japan each hold about one-third of the euro area's gravity. This structure has remained remarkably stable over time despite repeated shocks to the dollar's dominance.

JEL Classifications: F31, G12, G15, G28
Keywords: global FX funding; covered interest parity (CIP); funding gravity; financial power; U.S. dollar dominance.

*We gratefully acknowledge funding from the NBER-OFR Grant on Financial Frictions and Systemic Risk, the Jacobs Levy Equity Management Center for Quantitative Financial Research at the Wharton School, 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

Countries jostle for economic prowess through trade, innovation, and increasingly, through finance. Beyond output or productivity, financial strength shapes how nations influence global capital flows and currency use (e.g., Coppola, Krishnamurthy, and Xu, 2025; Clayton, Maggiori, and Schreger, 2025). Yet financial power is difficult to quantify.

Intuitively, a country's ability to attract capital when its own financial conditions shift—whether in debt, equity, or foreign exchange (FX) markets—reveals its structural dominance in global finance. We formalize this intuition in a multi-country game and develop a new measure of financial power, termed funding gravity. Funding gravity quantifies how much capital a country draws in response to changes in its funding conditions, and is a key consideration for countries contemplating interventions in own funding conditions. As such, funding gravity traces the evolution of financial dominance over time and highlights the investors central to its rise and fall.

We empirically study countries' funding gravity in the global FX funding market, a core venue for cross-border funding and investments. In this market, agents obtain short-term funding in one currency by pledging another as collateral through FX forwards and swaps. Because FX funding is instrumental in facilitating international capital flows, it is actively monitored and shaped by policymakers. During both the Global Financial Crisis of 2007-09 and the COVID-19 pandemic, for instance, the Federal Reserve intervened decisively to stabilize U.S. dollar (USD) FX funding conditions, including by establishing central bank swap lines.

Our estimates reveal a striking hierarchy of financial power. The funding gravity of the United States is as large as that of all other countries combined. The euro area commands roughly half the U.S. level, while the United Kingdom and Japan each account for about one-third of the euro area's gravity. This ordering has remained remarkably stable over time despite repeated shocks to the dollar's dominance, including trade tensions and global crises.

To obtain these results, we confront two central challenges. The first arises from the interconnected nature of global financial markets, which makes funding gravity inherently relative. Because investors allocate capital globally, a country's ability to attract funding depends not only on its own funding conditions but also on those of others. When one country adjusts its funding terms, others respond, creating endogeneity across the system. Further complicating the analysis, the amount of capital drawn to a country can also vary with underlying economic fundamentals, which may co-move across countries beyond the strategic interaction of funding conditions.

We address this issue by introducing funding blocs: ecosystems of countries that share

common funding needs and experience correlated shocks. As such, each funding bloc represents a distinct dimension of strategic interaction among economic agents and countries. Within a bloc, countries on the same side of the funding position cooperate, as an improvement in one country's funding conditions draws capital toward others in the same group, while countries on the opposite side compete for the same pool of global funding. A country may belong to several blocs simultaneously, much like overlapping alliances such as the Asia-Pacific Economic Cooperation (APEC) and the Transatlantic Trade and Investment Partnership (TTIP). Crucially, we construct these blocs so that their shocks are mutually orthogonal, ensuring that funding in one bloc responds only to conditions within that bloc and not to others. This orthogonality allows us to decompose a country's funding gravity into two components: (i) its participation across blocs and (ii) each bloc's intrinsic ability to attract capital. The approach not only yields a feasible empirical strategy but also illuminates the sources of financial dominance. The United States' exceptionally high funding gravity, for example, reflects its deep integration with the two most influential funding blocs rather than the breadth of its participation.

The second challenge is that capital flows in and out of the FX funding market are difficult to measure. Comprehensive data are scarce, especially at the global level, and even when positions are observed, they rarely align with the relevant funding horizons. We address this challenge with a novel dataset and an accompanying estimation methodology that recovers the funding quantities consistent with observed funding prices.

Our proprietary data cover outstanding FX forward and swap positions across eight major currencies and eighty investor groups from eight different geographies between 2012 and 2024. The dataset includes precise information on contract maturity, which allows us to estimate positions at specific funding horizons rather than at contract tenor alone. The price of FX funding is the spread between the FX-implied and direct risk-free borrowing rates, also known as deviations from covered interest parity (CIP). This price is not zero because marginal funding provider faces a budget constraint, the shadow cost of which varies with funding conditions (Du, Hébert, and Huber, 2023). The corresponding funding quantity must therefore include all positions that draw on the budget over a given period. For a three-month horizon, for example, we aggregate (i) all contracts with exactly three-month tenor, (ii) the full contribution of longer-dated contracts spanning that window, and (iii) the pro-rata share of shorter-dated contracts that overlap it.

The resulting measure of FX funding quantity reveals several patterns that validate our approach and shed light on the structure of the global FX market. First, roughly two-thirds of total funding positions at a given horizon arise from contracts with non-matching maturities, underscoring the importance of accounting for the maturity structure. Second,

FX forwards account for about one-quarter of the total funding obtained through swaps. Although forwards involve only a terminal exchange, counterparties typically hedge the initial FX risk with an offsetting spot trade, leaving the same net funding position as a swap. Hence, forwards, like swaps, contribute directly to FX-funding position (Dao, Gourinchas, and Itskhoki, 2025; De Leo, Keller, and Zou, 2025) and must be accounted. These findings confirm that our constructed quantity measure provides a comprehensive representation of FX funding activity and materially improves on approaches that count only same-tenor swaps. With both price and quantity measures in hand, we identify the independent funding blocs that underpin the global structure of financial power.

We construct funding blocs whose innovations in both funding quantities and prices are uncorrelated across blocs. This condition is equivalent to having independent shocks to funding conditions and underlying economic fundamentals, such that quantities in a bloc responds only to conditions in that bloc and not other blocs. Using the three-month horizon as our baseline, we find that the leading bloc is a Dollar bloc, in which agents obtain FX funding in USD against all other currencies, and vice versa. This bloc explains 74% of the variance in daily changes in investors' funding positions. The second bloc, a Pan-European bloc, reflects demand for the euro (EUR) and the pound sterling (GBP) FX funding, primarily against the yen (JPY) with a smaller USD component; it accounts for 10% of the variance. The explanatory power of the remaining blocs decays quickly and is small relative to the first two. The emergency of the Pan-European bloc is important as it shows that not all FX funding is driven by one Dollar funding cycle (Miranda-Agrippino and Rey, 2020). Equally important is what is missing: high-interest-rate currencies typically used in carry trades receive negligible weights in the two most prominent blocs. This pattern contrasts with An and Huber (2025), who find that in direct (unhedged) FX flows, carry trades form the second most important trading pattern, highlighting a potential divergence between the drivers of unhedged currency flows and exchange-rate-hedged FX funding.

Next, we quantify each country's funding gravity using the estimated funding blocs. Because countries participate in different blocs, a country's funding gravity depends on both the capital each bloc can attract and the intensity of the country's participation across blocs. We formalize a bloc's intrinsic ability to attract capital as proportional to the capital active within it; its "capital pull" is therefore proxied by the share of variance it explains in investors' daily funding-position changes. Since the blocs' capital pulls are orthogonal by construction, a country's overall funding gravity is simply the sum across blocs of each bloc's capital pull multiplied by the country's participation weight in that bloc. In equilibrium, a country's funding gravity directly captures its incentive to improve its funding condition.

Across all sample countries, the U.S. holds 50% of the total funding gravity, the euro

area 26%, the U.K. 9%, and Japan 8%. The remaining countries, Australia, Canada, New Zealand, Sweden, and Switzerland together comprise small residual shares. The dominance of the U.S. and the euro area reflects their high participation in the two leading funding blocs and underscores their clout in global financial power. Finally, this financial hierarchy has been remarkably persistent over time, reflecting the sustained capital pull of each funding bloc.

This paper contributes to three related strands of the literature: (i) currency dominance and global financial competition, (ii) geopolitics and financial power, and (iii) FX funding and CIP deviations. To start, currency dominance, exemplified by the U.S. dollar's century-long pre-eminence, endows a currency with a unique role in the international transmission of shocks (e.g., Miranda-Agrippino and Rey, 2020; An and Huber, 2025). Scholars have investigated how such dominance is reinforced through trade invoicing (e.g., Gopinath, Boz, Casas, Díez, Gourinchas, and Plagborg-Møller, 2020) and debt denomination (e.g., Coppola, Krishnamurthy, and Xu, 2025). This paper examines currencies' pre-eminence in the global FX funding market, which is critical to both trade and finance and has grown increasingly important since the 2008 crisis (Du and Huber, 2025). Our central result is the quantification of each country's funding gravity, revealing corresponding currencies' relative strength as they compete for investors' capital. We confirm the U.S. dollar's dominance, but also illustrate the role of other currencies; among others, we offer a new perspective on the international role of the euro (European Central Bank, 2025). In this vein, we complement the mostly theoretical literature on currency competition (e.g., Matsuyama, Kiyotaki, and Matsui, 1993; Zhang, 2014; Abadi, Fernàndez-Villaverde, and Sanches, 2025) by offering a direct quantification that illustrates the hierarchy in currency dominance.

Second, our analysis of financial power connects to the literature on geopolitics. There is a burgeoning theoretical literature on competitions among countries to amass and exert geopolitical power (e.g., Clayton, Maggiori, and Schreger, 2025; Broner, Martin, Meyer, and Trebesch, 2025). Progress has also been made to derive empirical proxies of geopolitical pressure (e.g., Clayton, Coppola, Maggiori, and Schreger, 2025, Bianchi, Horn, Rosso, and Sosa-Padilla, 2025). We share the common feature across these studies: that countries strategically compete and cooperate with each other for global influence. Our contribution is to recover these strategic dimensions empirically from revealed actions. Indeed, we use joint price and quantity data on FX funding to identify funding blocs that are characterized by mutually orthogonal shocks to funding prices and quantities. Each funding bloc thus represents distinct dimensions of strategic interaction among economic agents and countries, within which countries on the same side of the funding position cooperate, while countries on the opposite side compete for the same pool of global funding.

Lastly, we contribute to the literature on global currencies flows and CIP deviations. Early studies document the post-2008 breakdown of CIP and link it to arbitrage constraints faced by intermediaries (e.g., Du, Tepper, and Verdelhan, 2018; Cenedese, Della Corte, and Wang, 2021; Du, Hébert, and Huber, 2023). More recent work leverages granular quantity data to broaden the set of drivers behind CIP deviations. For example, recent papers highlight market frictions such as safe-asset scarcity and bank market power (Moskowitz, Ross, Ross, and Vasudevan (2024)), demand for repurchase agreement (Kloks, Mattille, and Ranaldo, 2024; Du, Strasser, and Verdelhan, 2025), and speculative demand for FX forwards (Dao, Gourinchas, and Itskhoki, 2025; De Leo, Keller, and Zou, 2025). Our conceptual insight is that even absent market frictions, CIP deviations can arise because the marginal funding provider faces a limited budget. Understand these deviations therefore requires a clean measure of the quantity that weighs on marginal agent's budget. Toward this objective, we make two empirical contributions. First, we construct an improved measure of the quantity relevant for CIP deviations at a given horizon by aggregating both forward and swap positions across maturities, including those that are not tenor-matched. Second, because funding quantities and CIP deviations are correlated across countries, we study their joint behavior through independent funding blocs, where funding quantity and price interact within blocs but not across. As such, these funding blocs capture distinct dimensions of variation that affect CIP deviations and weave the complex global funding market that we observe.

The next section outlines our baseline model of country competition and defines funding gravity and funding blocs. Section 3 describes the data. Section 4 constructs FX funding quantities from outstanding forwards and swaps, and Section 5 documents key facts about these quantities. Section 6 combines FX funding quantity with price data to construct and interpret funding blocs. Section 7 estimates countries' funding gravity and establishes currency hierarchy in global FX funding. Section 8 concludes.

2 Theoretical Foundation

We study a stylized FX funding environment in which real-money investors seek funding to acquire foreign assets, and intermediaries accommodate these flows. A country's funding condition captures intermediaries' marginal cost in that currency, and relative prices compare funding conditions across countries. Funding demand for a given country responds to a country-specific demand shifter as well as to the full set of funding conditions across countries. Each country chooses its funding condition to maximize the foreign funding it attracts, with aggregate net demand summing to zero across countries by construction. To

highlight the central mechanism, the baseline assumes a perfectly elastic market: funding prices are determined by funding conditions and do not move with the level of aggregate demand. Investor heterogeneity is collapsed into a single observable net demand per country aggregated from all other countries. This baseline isolates each country's funding choice; later we discuss how to incorporate inelastic market and investor heterogeneity.

There are $N+1$ countries for $n = 0, \ldots, N$, with country 0 serving as the numeraire. Each country n chooses the funding condition $f_{n,t}$ in each period t, values the net aggregate demand from all other countries $Q_{n,t}$, c is the quadratic adjustment cost, and $\tilde{f}_{n,t}$ is the underlying funding condition that exogenously moves around:

$$\max_{f_{n,t}} \; Q_{n,t} - \frac{c}{2}\big(f_{n,t} - \tilde{f}_{n,t}\big)^2. \tag{1}$$

The relative price of demand funding for country n from country m is given by the relative funding conditions,

$$P_{n,m,t} = f_{m,t} - f_{n,t}, \tag{2}$$

so better funding conditions make funding cheaper.

Because investors can easily substitute across countries, we model the aggregate net demand for country n from all other countries as

$$Q_{n,t} = a_{n,t} + g_n\big(f_{0,t}, f_{1,t}, \ldots, f_{N,t}\big), \tag{3}$$

where $a_{n,t}$ is a country-n specific demand shock that satisfies $\sum_{n=0}^{N} a_{n,t} = 0$ and $g_n(\cdot)$ is a country-specific function capturing how investor demand responds to the funding conditions of all countries.

To model $g_n(\cdot)$, we consider N segments, where the k-th segment is given by a vector of given weights $\{w_{n,k}\}$ satisfying $\sum_{n=0}^{N} w_{n,k} = 0$. Its price is given by the relative funding condition:

$$P^{\text{seg}}_{k,t} = -\sum_{n=0}^{N} w_{n,k} f_{n,t} = \sum_{n=1}^{N} w_{n,k} P_{n,0,t}. \tag{4}$$

A lower price for $P^{\text{seg}}_{k,t}$ attracts more funding to segment k, which is assumed to be proportional to the amount of capital μ_k allocated to segment k,

$$-\theta \mu_k P^{\text{seg}}_{k,t}, \tag{5}$$

where θ is a sensitivity parameter. These segment demands are then allocated to country n

based on the weight $w_{n,k}$, so the function $g(\cdot)$ is

$$g_n\big(f_{0,t}, f_{1,t}, \ldots, f_{N,t}\big) = -\sum_{k=1}^{N} w_{n,k}\theta\mu_k P_{k,t}^{\text{seg}}. \tag{6}$$

The following proposition solves the equilibrium funding conditions that each country optimally chooses. Appendix A.1 provides a proof.

PROPOSITION 1. *There exists a unique equilibrium. The equilibrium funding condition of country n is*

$$f_{n,t} = \tilde{f}_{n,t} + \frac{\theta\varphi_n}{c}, \tag{7}$$

where

$$\varphi_n := \sum_{k=1}^{N} \mu_k w_{n,k}^2. \tag{8}$$

The key parameter φ_n captures the incentive of country n to improve its funding condition. It reflects how strongly a country is tied to the global funding system through its participation in different funding segments. Specifically, φ_n depends on the assortative matching between a country's relative importance in each segment, represented by the weight $w_{n,k}$, and the amount of capital μ_k active in that segment. Intuitively, a country that plays a central role in a large and liquid funding segment has a greater incentive to improve its funding condition, since doing so affects a larger share of global capital. For example, as will be shown later, the United States has a high φ_n because the USD dominates the most traded segment (the dollar segment), giving it the strongest incentive to improve its funding condition. In contrast, a smaller country such as Australia, whose currency primarily participates in a less traded segment, faces a much smaller incentive to do so.

Using Δ to denote the period-by-period change from $t-1$ to t, we have from equations (4) and (7),

$$\Delta P_{k,t}^{\text{seg}} = \sum_{n=1}^{N} w_{n,k}\Delta P_{n,0,t} = \sum_{n=1}^{N} w_{n,k}(\Delta\tilde{f}_{0,t} - \Delta\tilde{f}_{n,t}). \tag{9}$$

The net demand $Q_{k,t}^{\text{seg}}$ for each segment k is solved from the equation

$$Q_{n,t} = \sum_{k=1}^{N} w_{n,k} Q_{k,t}^{\text{seg}}. \tag{10}$$

To explicitly express $Q_{k,t}^{\text{seg}}$, denote the $N \times N$ matrix $\mathbf{W} := \{w_{n,k}\}_{n,k=1,\ldots,N}$, and the (k,n)-th

element of $\mathbf{W}^{-1}$ as $v_{k,n}$. Then, equations (3) and (6) implies

$$\Delta Q^{\text{seg}}_{k,t} = \sum_{n=1}^{N} v_{k,n} \Delta Q_{n,t} = \sum_{n=1}^{N} v_{k,n} \Delta a_{n,t} - \theta \mu_k \Delta P^{\text{seg}}_{k,t}. \tag{11}$$

Equations (9) and (11) show that changes in the equilibrium price ΔP and demand ΔQ are driven by underlying shocks to funding conditions $\Delta \tilde{f}$ and attractiveness Δa. These underlying shocks can be correlated across countries, which in turn generate correlated observations of $\Delta Q_{n,t}$ and $\Delta P_{n,t}$ for individual countries. The next proposition demonstrates that, based on country-level observations, one can select the weights $\{w_{n,k}\}$ to construct funding segments such that both the equilibrium price changes ΔP and demand changes ΔQ are mutually uncorrelated across segments. Appendix A.2 provides a proof.

PROPOSITION 2. *The weights $\{w_{n,k}\}$ can be chosen so that different segments satisfy:*

1. *Uncorrelated price changes:*

$$\operatorname{cov}\left(\Delta P^{\text{seg}}_{k,t}, \Delta P^{\text{seg}}_{j,t}\right) = 0 \quad \textit{for any } k \neq j. \tag{12}$$

2. *Uncorrelated demand changes:*

$$\operatorname{cov}\left(\Delta Q^{\text{seg}}_{k,t}, \Delta Q^{\text{seg}}_{j,t}\right) = 0 \quad \textit{for any } k \neq j. \tag{13}$$

To ensure that these segments also capture independent shocks to funding conditions $\Delta \tilde{f}$ and attractiveness Δa, we assume that the underlying attractiveness shock for one segment is uncorrelated with the underlying funding condition shock of another. Specifically, for any pair of distinct segments $k \neq l$,

$$\operatorname{cov}\left(\sum_{n=1}^{N} v_{k,n}\, \Delta a_{n,t}, \sum_{n=1}^{N} w_{n,l} \left(\Delta \tilde{f}_{0,t} - \Delta \tilde{f}_{n,t}\right)\right) = 0. \tag{14}$$

An immediate corollary is that the segments identified in Proposition 2 also exhibit mutually uncorrelated underlying shocks to funding conditions $\Delta \tilde{f}$ and attractiveness Δa. In other words, these segments recover independent structural shocks to global funding markets.

COROLLARY 1. *The segments exhibit mutually uncorrelated price changes (12) and uncorrelated demand changes (13) if and only if they satisfy:*

1. *Uncorrelated funding shocks:*

$$\operatorname{cov}\left(\sum_{n=1}^{N} w_{n,k}\left(\Delta\tilde{f}_{0,t} - \Delta\tilde{f}_{n,t}\right), \sum_{n=1}^{N} w_{n,j}\left(\Delta\tilde{f}_{0,t} - \Delta\tilde{f}_{n,t}\right)\right) = 0 \quad \textit{for any } k \neq j. \tag{15}$$

2. *Uncorrelated attractiveness shocks:*

$$\operatorname{cov}\left(\sum_{n=1}^{N} v_{k,n}\Delta a_{n,t}, \sum_{n=1}^{N} v_{j,n}\Delta a_{n,t}\right) = 0 \quad \textit{for any } k \neq j. \tag{16}$$

Empirically measuring $\varphi_n = \sum_{k=1}^{N} \mu_k w_{n,k}^2$ requires both the weights $w_{n,k}$ and the amount of active capital μ_k in each segment k, where the latter is not directly observable. We assume that μ_k is proportional to the variance of demand changes, $\operatorname{var}\left(\Delta Q_{k,t}^{\text{seg}}\right)$. Specifically, we set $\mu_k = L \times \operatorname{var}\left(\Delta Q_{k,t}^{\text{seg}}\right)$ for all $k = 1, \ldots, N$ and some proportionality constant L. The intuition is that a larger $\operatorname{var}\left(\Delta Q_{k,t}^{\text{seg}}\right)$ reflects more active trading in segment k, implying that more capital is deployed in that segment. Since the scaling constant L is unobserved and does not affect relative magnitudes, we set $L = 1$ without loss of generality. The empirical measure of φ_n is therefore given by

$$\varphi_n = \sum_{k=1}^{N} \operatorname{var}\left(\Delta Q_{k,t}^{\text{seg}}\right) w_{n,k}^2. \tag{17}$$

The baseline model assumes a perfectly elastic market, meaning that funding prices do not adjust with the level of aggregate demand. Investor heterogeneity is also abstracted away, represented by a single observable net demand per country aggregated from all other countries. Appendix B extends the framework to incorporate market inelasticity and heterogeneous investors. We show that, to obtain uncorrelated funding shocks and uncorrelated attractiveness shocks as in Proposition 2 and Corollary 1, it is equivalent to jointly orthogonalize equilibrium price changes and quantity changes aggregated across investors, a procedure that we implement empirically.

3 Data

We outline the data used to measure prices and quantities in global FX funding markets. Throughout, the U.S. dollar (USD) is the numeraire: exchange rates are expressed as number of foreign currencies per USD, and quantities are measured in USD.

3.1 Price Data

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

$$R^{\text{syn}}_{n,t,\tau} \;:=\; R_{\$,t,\tau}\left(\frac{F_{n,t,t+\tau}}{S_{n,t}}\right)^{12/\tau}, \tag{18}$$

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:

$$P_{n,t,\tau} \;:=\; \ln R^{\text{syn}}_{n,t,\tau} - \ln R_{n,t,\tau}. \tag{19}$$

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).

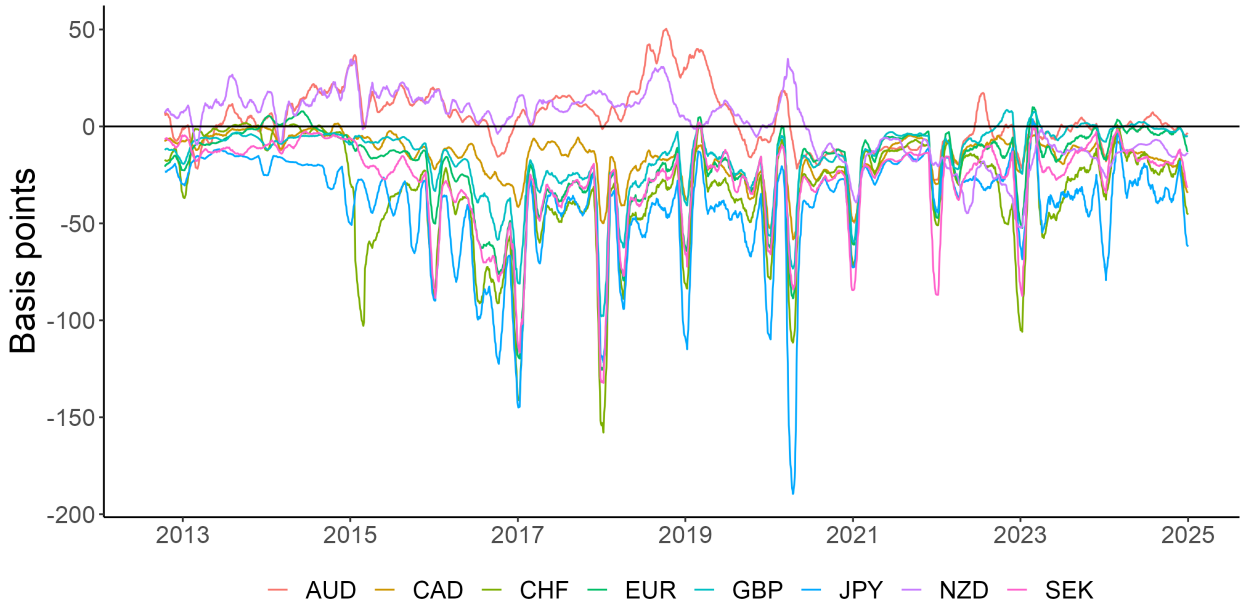
Similarly, for any two non-USD currencies n and n', the price of demanding currency n funding and supplying currency n' funding is defined by the triangular arbitrage:

$$P_{n,n',t,\tau} = P_{n,t,\tau} - P_{n',t,\tau}. \tag{20}$$

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, whereas LIBOR embeds bank credit/liquidity risk (also discontinued around 2021) and Treasury yields can be distorted by specialness and regulatory demand. OIS quotes are available for the USD and eight non-USD currencies: Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), euro (EUR), British pound (GBP), Japanese yen (JPY), New Zealand dollar (NZD), and Swedish krona (SEK).

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 subpe-

Figure 1: Price of FX funding across currencies ($P_{n,t,1\mathrm{M}}$) at 1-month tenor



Notes: This figure plots $P_{n,t,1\mathrm{M}}$, the 1M 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.

riods. 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. To span short- and long-tenor funding, we study 1-month (1M), 3-month (3M), 6-month (6M), and 1-year (1Y) tenors. Figure 1 plots $P_{n,t,1\mathrm{M}}$, the 1M price of demand funding in currency n by supplying USD funding for each non-USD currency n, as defined in (19). To reduce high-frequency noise, the series are shown as 30-day moving averages. 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.

3.2 Quantity Data

Our quantity data come from the CLS Group (CLS), which settles FX trades for 72 settlement members, primarily large multinational banks.[1] As the largest single source of FX execution data, CLS covers over 50% of global FX volumes.

We use CLS daily aggregates of outstanding FX swap and forward positions; Section 3.3

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

shows why including forwards, rather than swaps alone, is essential. Each trading day, CLS reports end-of-day gross outstanding notional buys and sells by investor type and location (details on investor taxonomy below). Concretely, for each investor-type-location $\times$ currency pair $\times$ maturity bucket $\times$ trading day, we observe (i) the total outstanding notional on the buy and sell sides against all counterparties (including those in the same investor category) and (ii) the outstanding-amount-weighted average remaining maturity within that bucket.[2]

Tracking remaining maturity is essential for mapping outstanding quantities to direct exposure to FX funding for a given funding horizon. For example, a six-month EUR borrowing via USD initiated at date t rolls down to a three-month exposure after three months and should load on the three-month rather than the six-month FX-funding price. In Section 4 we show that this roll-down of existing positions accounts for a sizable share of the total exposures.

The raw CLS panel spans 17 currencies from September 2012 to December 2024; for consistency with our price benchmarks, we restrict the analysis to the nine currencies with OIS curves.[3] We convert pair-level positions into currency-level funding demand and supply. Specifically, any position that borrows currency n against any currency n' is classified as demand for funding in n; the opposite direction is supply of n'. For example, positions in EUR/JPY or EUR/USD that borrow EUR are counted as demand for EUR funding; the offsetting legs supply JPY or USD, respectively. All quantities are expressed in USD billions using contemporaneous spot rates, consistent with our USD-numeraire convention.

We classify investors into three groups: (i) banks that settle with CLS, (ii) banks that do not settle with CLS, and (iii) customers. For CLS-settling banks, we observe both the physical location and the headquarter country, so we index groups by $Bank_{Loc}^{HQ}$, where $Loc, HQ \in \{US, EZ, UK, JP, AU, CH, CA, ROW\}$ corresponding to United States (US), Euro Zone (EZ), United Kingdom (UK), Japan (JP), Australia (AU), Switzerland (CH), Canada (CA), and Rest of World (ROW). For example, $Bank_{US}^{UK:HQ}$ denotes banks located in the United States but headquartered in the United Kingdom. For non-CLS-settling banks, headquarters is not observed, so we write $Bank_{Loc}^{other}$ (e.g., $Bank_{US}^{other}$). For customers, only location is observed, and we write $Customer_{Loc}$ (e.g., $Customer_{US}$).

[2]CLS defines remaining maturity as the number of calendar days from the current date to the settlement date of the far leg, which is typically two business days after the contractual far date. The maturity buckets are: tomorrow-next and spot-next (overnight variants with remaining maturity ≤ 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 ($\geq$ 369 days).

[3]The full CLS currency list comprises the nine OIS currencies plus Danish krone (DKK), Hong Kong dollar (HKD), Israeli new shekel (ILS), Korean won (KRW), Mexican peso (MXN), Norwegian krone (NOK), Singapore dollar (SGD), and South African rand (ZAR).

3.3 How Swaps and Forwards Contribute to FX Funding

This subsection explains why both FX swaps and outright forwards belong in our measure of FX funding, and measures their relative importance.

An FX swap consists of a near-date exchange of the two currencies at the spot rate together with a reverse exchange at $t+\tau$ at a forward rate fixed today. An outright forward has only the far leg: the two currencies are exchanged at $t+\tau$ at a rate agreed at t. Hence, an outright forward can be written as an FX swap plus an offsetting spot trade. The offsetting spot leg carries no funding over the horizon; the swap leg embodies the borrowing/lending between the two currencies. Therefore, a forward contributes to FX funding one-for-one with an otherwise identical swap of the same direction and horizon.

One may argue that end-users may transact forwards to hedge or speculate rather than to obtain funding per se (e.g., exchange-rate appreciation/depreciation views; see Shin, Wooldridge, and Xia, 2025). However, the counterparty that intermediates the forward typically hedges FX risk with a spot trade, stripping out exchange-rate risk and leaving the same funding exposure as a swap. Consequently, from the perspective of direct balance-sheet exposures, forwards demand the same funding capacity as swaps.

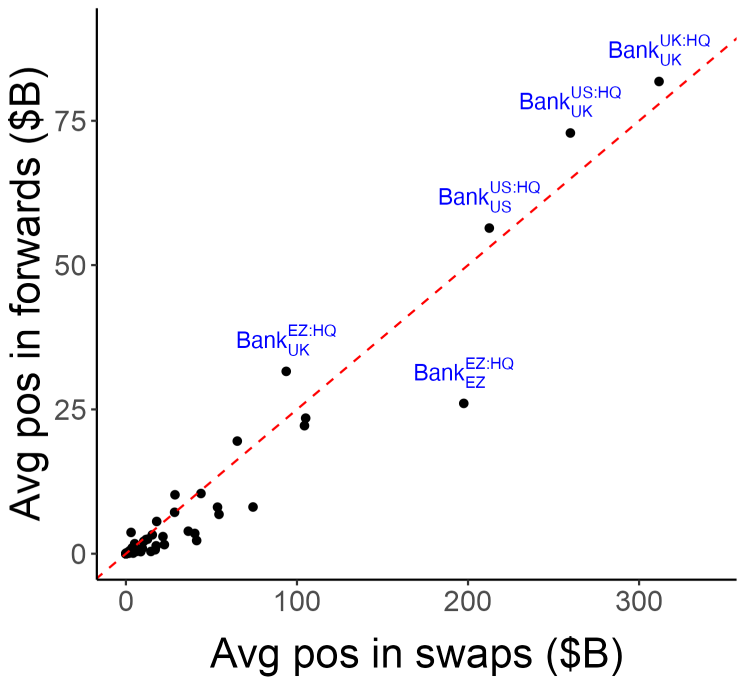
To gauge the quantitative importance of swaps versus forwards, Figure 2 plots, for each investor-type-location group, the time-series average of end-of-day outstanding gross notionals in swaps (x-axis) and forwards (y-axis) at the 1M tenor. Results for other tenors are similar and omitted for brevity. Two facts emerge. First, larger investor groups use both instruments more, with points lying close to a line through the origin, showing that forward usage scales roughly proportionally with swap usage. Second, on average, about 80% of FX funding is executed via swaps and about 20% via forwards. Forwards therefore account for a material share of FX funding and should not be ignored by analyses that focus solely on swaps.

4 Constructing Quantities in FX Funding

In this section, we map the raw outstanding positions into direct exposures to FX funding and document their time-series properties.

Our objective is to recover the exposures 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. Thus, even if an investor has no outstanding 1-month position today, a previously entered 3-month USD borrowing still spans the next month and contributes to the 1-month horizon during that window. Accordingly, our direct exposures over a funding horizon τ aggregates outstanding contracts across all remaining maturities, weighting each

Figure 2: Swaps versus forwards by investor group and tenor



Notes: Each dot represents an investor-type-location group. The axes display the time-series mean of end-of-day outstanding gross notionals, with swaps on the x-axis and forwards on the y-axis, averaged across all currencies at the 1M tenor. Labeled points identify the largest groups by average notional.

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 the mismeasurement that would result from looking only at contracts with tenor τ.

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

$$\mathcal{E}^D_{i,n,t,\tau} := \sum_{m\in\{1\text{W},1\text{M},3\text{M},6\text{M},1\text{Y},\text{Longer}\}} O^D_{i,n,t,m}\,\mu^D_{i,n,t,m}(\tau), \quad \mu^D_{i,n,t,m}(\tau) := \frac{\min\{\tau,\, M^D_{i,n,t,m}\}}{\tau} \in [0,1], \tag{21}$$

where m indexes the maturity buckets, $O^D_{i,n,t,m}$ is the end-of-day outstanding total notional on the demand side in bucket m, and $M^D_{i,n,t,m}$ is the corresponding outstanding-amount-weighted remaining maturity. The weight $\mu^D_{i,n,t,m}(\tau)$ captures the share of each position's remaining life that overlaps the horizon $[t, t+\tau]$: it equals 1 when the remaining maturity exceeds τ and scales linearly otherwise. Intuitively, a 1-month contract contributes 1/12 to the 1-year direct exposure, whereas a 1-year contract contributes fully (weight 1) to the 1-month direct exposure. We construct the supply side exposure analogously, $\mathcal{E}^S_{i,n,t,\tau} =$

Table 1: Contribution of maturity buckets m to horizon-τ direct exposure

Maturity bracket	1M Funding	3M Funding	6M Funding	1Y Funding
1W	2%	1%	1%	1%
1M	26%	14%	11%	9%
3M	45%	44%	35%	29%
6M	14%	22%	24%	21%
1Y	9%	14%	21%	26%
longer	3%	5%	8%	14%

Notes: Rows are CLS maturity buckets $\{1\text{W}, 1\text{M}, 3\text{M}, 6\text{M}, 1\text{Y}, \text{Longer}\}$; columns are funding horizon $\{1\text{M}, 3\text{M}, 6\text{M}, 1\text{Y}\}$. Each cell is the market-weighted share (percent) of horizon-τ exposure accounted for by bucket m, aggregated over all investor groups i, currencies n, and dates t.

$\sum_m O^S_{i,n,t,m}\, \mu^S_{i,n,t,m}(\tau)$; the net exposures are $\mathcal{E}_{i,n,t,\tau} = \mathcal{E}^D_{i,n,t,\tau} - \mathcal{E}^S_{i,n,t,\tau}$.

To quantify how maturity bucket m contributes to horizon-τ direct exposure, define the market-weighted share on the demand side:

$$\text{Share}_m(\tau) := \frac{\sum_{i,n,t} O^D_{i,n,t,m}\, \mu^D_{i,n,t,m}(\tau)}{\sum_{i,n,t} \mathcal{E}^D_{i,n,t,\tau}}. \tag{22}$$

The numerator aggregates, across investors i, currencies n, and dates t, the portion of horizon-τ demand exposure attributable to contracts in maturity bucket m; the denominator is total horizon-τ demand exposure. By market clearing, the corresponding supply-side shares are the same, so we report demand-based shares only. By construction, $\sum_m \text{Share}_m(\tau) = 1$.

Table 1 reports the market-weighted shares $\text{Share}_m(\tau)$ by maturity bucket m (rows) and funding horizon τ (columns). Two facts stand out. First, diagonal entries ($m = \tau$) are small—26% (1M), 44% (3M), 24% (6M), and 26% (1Y). Second, off-diagonal contributions are large; for example, the 3M bucket accounts for 45% of the 1M direct exposure. Thus, using only same-tenor contracts (e.g., 1M positions for 1M funding) misses more than half of the relevant quantity. Correct measurement requires the roll-down mapping in (21), which translates outstanding positions into horizon-τ exposures.

Next, we study what drives day-to-day changes in net direct exposure $\mathcal{E}_{i,n,t,\tau}$. These changes have two components: (i) the *aging* of existing positions as remaining maturities shorten, and (ii) *new* positions initiated between t and $t+1$ (the component typically observed in transaction-flow data). Empirically, new trading (ii) explains a substantial share of daily variation but not all—the remainder reflects the impact of aging. Therefore, the quantity relevant for horizon-τ pricing must track both components via our measure $\mathcal{E}_{i,n,t,\tau}$, rather

than relying solely on transaction flows for new positions.

Define the counterfactual next-day net exposure if no new trades occur and all positions simply age by one day as

$$\widehat{\mathcal{E}}_{i,n,t+1|t,\tau} \;:=\; \sum_m O^D_{i,n,t,m}\,\mu^{D,\text{age}}_{i,n,t,m}(\tau) \;-\; \sum_m O^S_{i,n,t,m}\,\mu^{S,\text{age}}_{i,n,t,m}(\tau), \tag{23}$$

with aging weights

$$\mu^{D,\text{age}}_{i,n,t,m}(\tau) = \frac{\min\{\tau,\,(M^D_{i,n,t,m}-1)_+\}}{\tau}, \qquad \mu^{S,\text{age}}_{i,n,t,m}(\tau) = \frac{\min\{\tau,\,(M^S_{i,n,t,m}-1)_+\}}{\tau}, \tag{24}$$

where $(x)_+ = \max\{x, 0\}$ ensures positions that mature between t and $t+1$ contribute zero at $t+1$.

The aging and new components are then

$$\Delta\mathcal{E}^{\text{age}}_{i,n,t+1,\tau} \;:=\; \widehat{\mathcal{E}}_{i,n,t+1|t,\tau} - \mathcal{E}_{i,n,t,\tau}, \tag{25}$$

$$\Delta\mathcal{E}^{\text{new}}_{i,n,t+1,\tau} \;:=\; \mathcal{E}_{i,n,t+1,\tau} - \widehat{\mathcal{E}}_{i,n,t+1|t,\tau}, \tag{26}$$

so that the identity for day-to-day changes in net exposures

$$\mathcal{E}_{i,n,t+1,\tau} - \mathcal{E}_{i,n,t,\tau} = \Delta\mathcal{E}^{\text{net,age}}_{i,n,t+1,\tau} + \Delta\mathcal{E}^{\text{net,new}}_{i,n,t+1,\tau} \tag{27}$$

holds exactly by construction. We then estimate a pooled regression across all (i, n, t, τ):

$$\Delta\mathcal{E}^{\text{new}}_{i,n,t+1,\tau} = \alpha + \beta(\mathcal{E}_{i,n,t+1,\tau} - \mathcal{E}_{i,n,t,\tau}) + \epsilon_{i,n,t,\tau}. \tag{28}$$

The estimate is $\beta = 0.82$, implying that about 82% of the day-to-day change in net exposure is driven by new trades, while the remaining 18% reflects the aging of existing positions. Thus, relying only on transaction flows would miss roughly one-fifth of the relevant variation in direct exposures.

5 Facts about FX Funding

This section documents stylized facts about investors' direct exposure in FX funding. We identify the largest demanders and suppliers by currency and funding horizon, and show that activity is highly concentrated: most investors demand and supply FX funding in only a small number of currencies.

Figure 3 summarizes who supplies and demands FX funding at the 1-month tenor. For

each currency $n \in \{\text{USD}, \text{EUR}, \text{JPY}\}$ and fixed $\tau = 1M$, we compute investor-group shares from the balance-sheet measures $\mathcal{E}^D_{i,n,t,\tau}$ and $\mathcal{E}^S_{i,n,t,\tau}$, average those shares over time t, identify the top three demanders and top three suppliers, and then plot their daily shares over the sample while keeping the investor identity fixed (shown in the legends). Three patterns emerge. (i) Concentration: in each currency, a small number of investor groups account for the bulk of supply and demand. (ii) Persistence: the composition of leaders is relatively stable over time. (iii) Geography: USD demand is led by Euro-area customers, EUR demand by U.S. customers, and JPY demand by Euro-area banks; on the supply side, USD supply is dominated by U.S. customers with notable contributions from UK banks, EUR supply by Euro-area customers and banks, and JPY supply by Japanese banks and Rest-of-World customers. Results for 3M, 6M, and 1Y tenors are similar and omitted for brevity.

To quantify how concentrated an investor's activity is across currencies, define the top-currency share for demand and supply. For investor i on date t and funding horizon τ, let

$$\text{TopDemShare}_{i,t,\tau} := \frac{\mathcal{E}^D_{i,n^D(i,t,\tau),t,\tau}}{\sum_n \mathcal{E}^D_{i,n,t,\tau}}, \qquad n^D(i,t,\tau) \in \arg\max_n \mathcal{E}^D_{i,n,t,\tau}, \tag{29}$$

and define $\text{TopSupShare}_{i,t,\tau}$ analogously using $\mathcal{E}^S_{i,n,t,\tau}$. We then average these shares across investors and time, separately for banks and customers, for each funding horizon $\tau \in \{1\text{M}, 3\text{M}, 6\text{M}, 1\text{Y}\}$. Figure 4 plots the results.

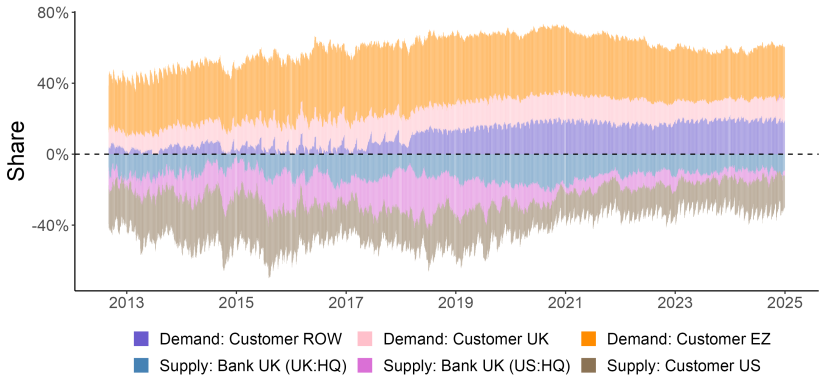
Two patterns emerge. First, activity is highly concentrated: a single currency typically accounts for a majority of an investor's funding demand and supply at every horizon. Second, concentration is systematically higher on the supply side than on the demand side, and higher for customers than for banks. Concentration is fairly stable across horizons, with at most a mild increase at longer horizons. These facts align with a preferred-habitat view, where investors demand and supply FX funding via their preferred currencies.

Having measured how concentrated activity is, we next report which currency most often anchors investors' demand and supply. Using the same $n^D(i,t,\tau)$ and $n^S(i,t,\tau)$ defined above, we compute the fraction of investor-days on which each currency is the most demanded/supplied currency, again separately for banks and customers and for each funding horizon.

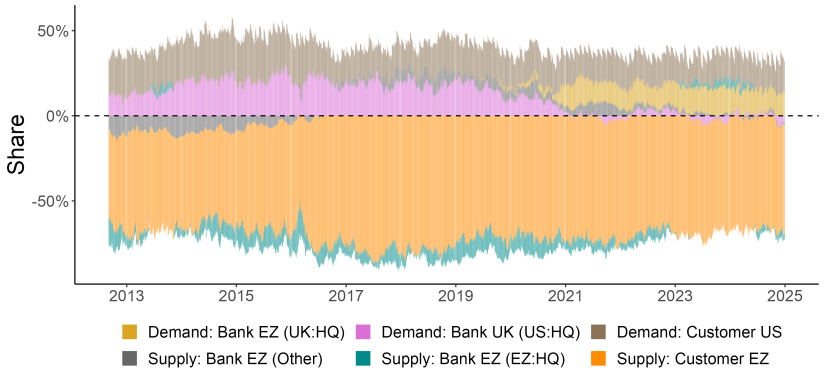
Figure 5 displays these distributions and reveals a clear asymmetry between investor types. On the demand side (top panel), banks are diversified across several currencies, whereas customers most often demand USD at every funding horizon. On the supply side (bottom panel), the pattern is largely the reverse: banks most frequently supply USD, while customers' modal supply currency is more often non-USD. The mirror-image structure points to the centrality of the dollar in global FX funding. USD dominates either as the final funding

Figure 3: Top 3 funding demanders and suppliers by currency

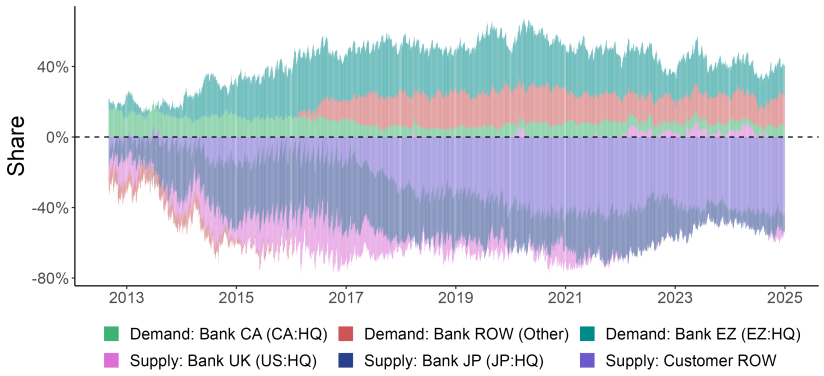
(a) USD



(b) EUR



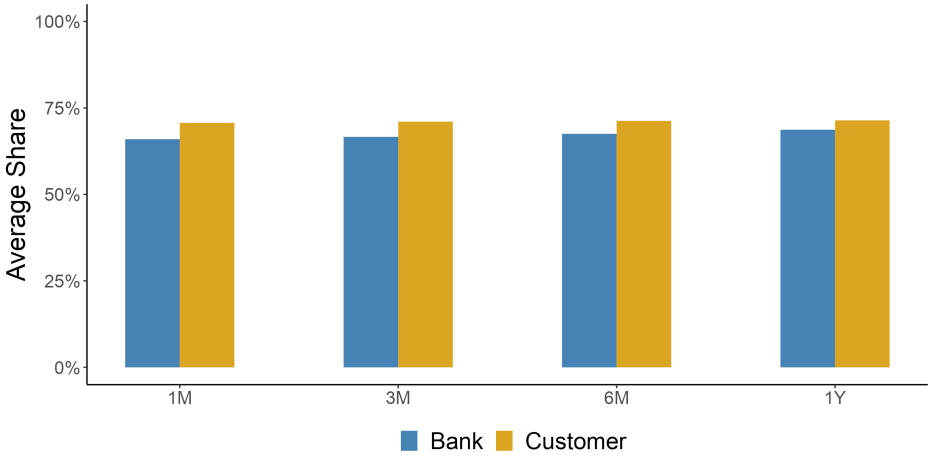
(c) JPY



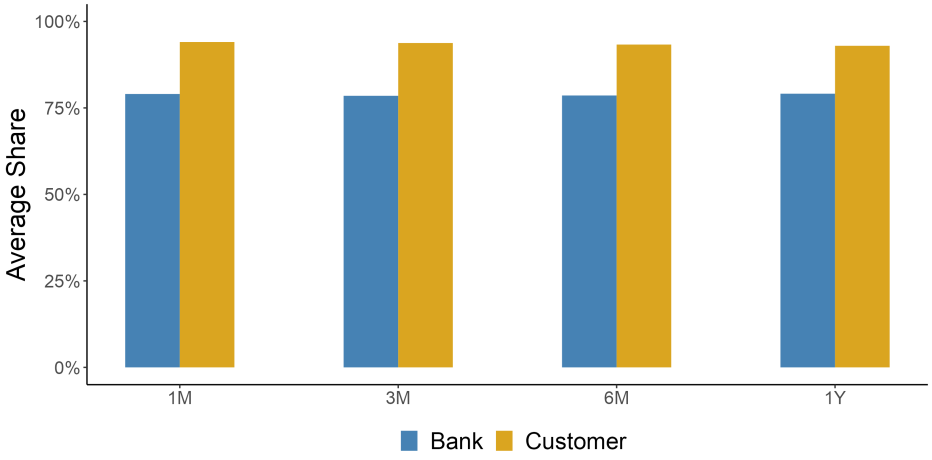
Notes: For each currency n with $\tau = 1\text{M}$, we compute investor-group shares from $\mathcal{E}^{D}_{i,n,t,\tau}$ (demand against all other currencies) and $\mathcal{E}^{S}_{i,n,t,\tau}$ (supply against all other currencies), rank groups by their time-averaged shares, fix the top three on each side, and plot their daily shares (positive for demand, negative for supply). Each group is represented by the same color across the panels.

Figure 4: Concentration in the top currency by funding horizon and investor type

(a) Demand



(b) Supply



Notes: $\text{TopDemShare}_{i,t,\tau}$ ($\text{TopSupShare}_{i,t,\tau}$) is the share of investor i's total demand (supply) at horizon τ accounted for by the single currency with the largest $\mathcal{E}^{D}_{i,n,t,\tau}$ ($\mathcal{E}^{S}_{i,n,t,\tau}$). Bars report the averages of these shares across investors and time, separately for banks and customers, for each funding horizon $\tau \in \{1\text{M}, 3\text{M}, 6\text{M}, 1\text{Y}\}$.

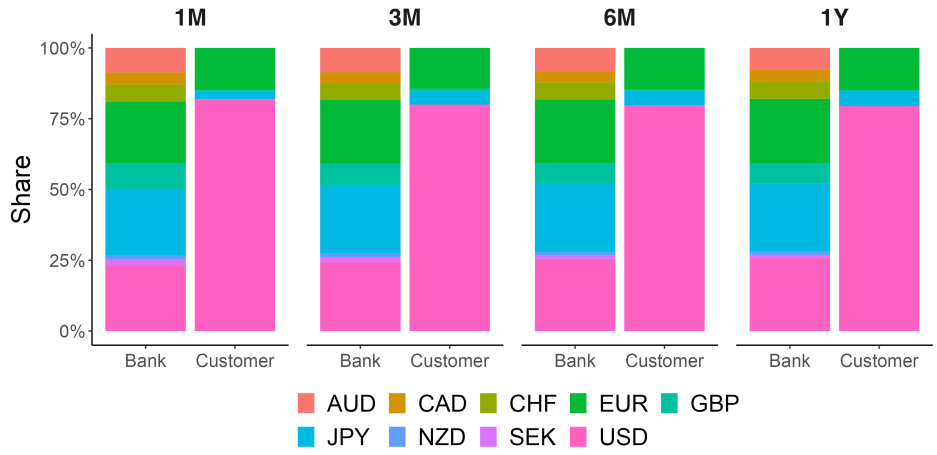
currency or as the bridge/vehicle currency to intermediate between two non-USD currencies.

6 Funding Blocs in FX

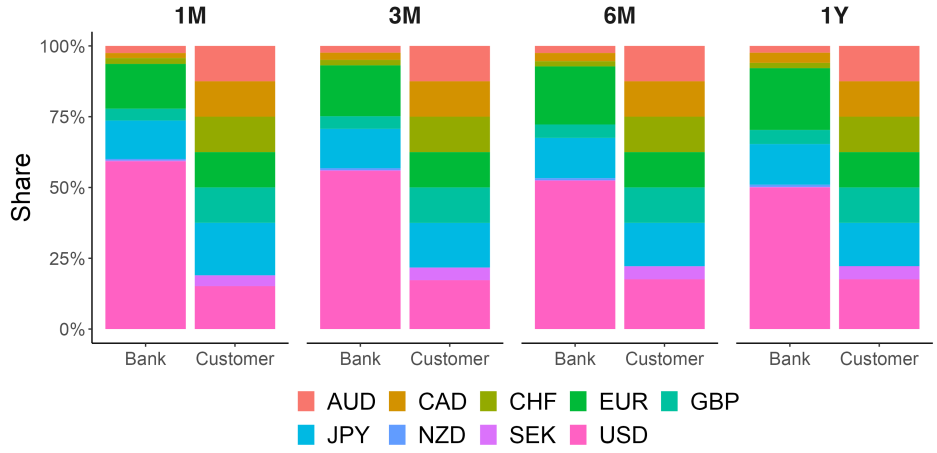
In this section we construct and interpret the blocs of global FX funding. We show that the global FX funding universe is well summarized by two dominant, mutually independent segments: (i) a Dollar segment (USD funding vis-à-vis all others) and (ii) a Pan-European

Figure 5: Most demanded and most supplied currency by funding horizon and investor type

(a) Most demanded currency



(b) Most supplied currency



Notes: For each investor i and date t at horizon τ, $n^D(i,t,\tau)$ ($n^S(i,t,\tau)$) is the currency with the largest $\mathcal{E}^D_{i,n,t,\tau}$ ($\mathcal{E}^S_{i,n,t,\tau}$). Bars show, for each investor type and funding horizon, the share of investor-days on which a given currency is the most demanded/supplied currency.

segment (EUR and GBP funding vis-à-vis all others). We then analyze these funding blocs. Three insights emerge. First, USD-based CIP deviations are in fact driven by movements in funding conditions in both funding segments. Second, large global banks supply substantial dollar funding but diversify by splitting it across the two segments. Third, the relative importance of Pan-European grows at longer horizons.

6.1 Empirical Findings

Table 2 reports the results of our procedure: the currency weights $w_{n,k}$ for each funding segment and the share of total exposure explained. We focus on 1M tenor here; results for

Table 2: Independent funding segments

	Seg 1	Seg 2	Seg 3	Seg 4	Seg 5	Seg 6	Seg 7	Seg 8
AUD	-0.01	-0.02	0.13	1.00	0.06	-0.03	-0.24	-0.74
CAD	-0.00	-0.01	0.04	0.04	-0.00	1.00	0.42	0.16
CHF	-0.05	0.02	-0.02	-0.02	1.00	0.05	-0.25	0.10
EUR	-0.56	1.00	-0.70	0.34	-0.88	0.07	-0.32	0.19
GBP	-0.17	0.43	1.00	-0.57	-0.15	-0.25	-0.49	-0.16
JPY	-0.18	-0.94	0.03	-0.07	-0.27	0.00	0.00	0.06
NZD	-0.01	-0.02	0.02	0.07	0.01	-0.03	0.11	1.00
SEK	-0.01	-0.01	0.01	0.02	0.03	-0.07	1.00	-0.14
USD	1.00	-0.45	-0.52	-0.81	0.19	-0.74	-0.25	-0.46
Var explained	83.0%	10.0%	3.0%	2.0%	1.0%	0.0%	0.0%	0.0%

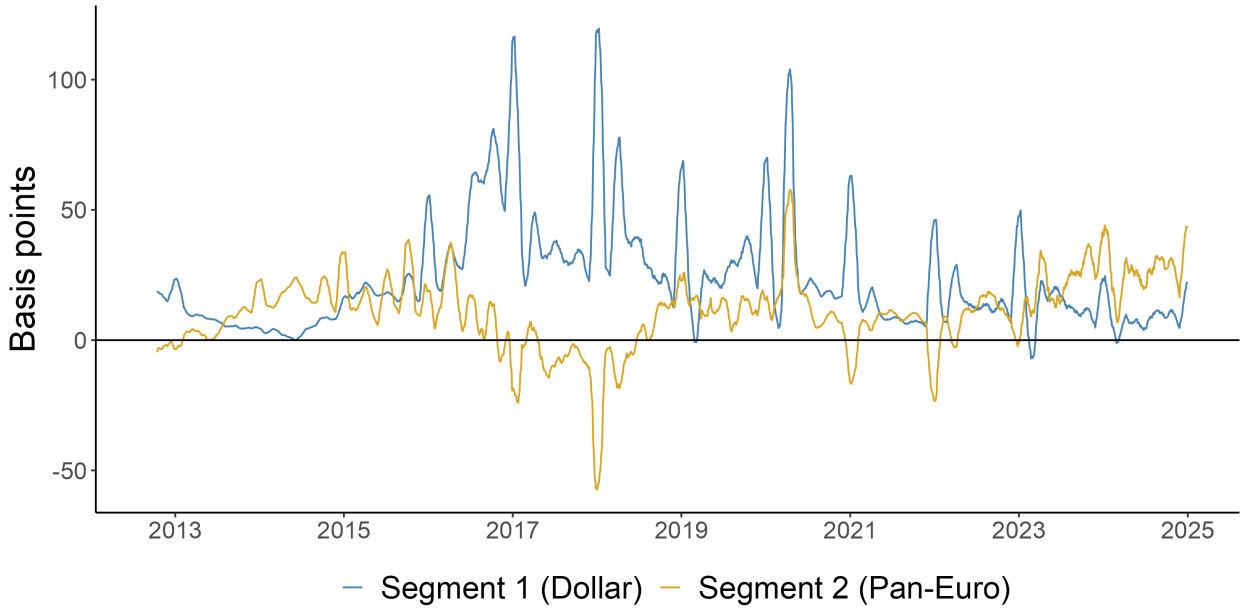
Notes: Currency weights are constructed to satisfy (A25) and (A26) and are sign/scale-normalized for interpretability (Segment 1: USD set to +1; Segment 2: EUR set to +1). "Var explained" is the share of total exposures explained by each segment.

other tenors are similar and omitted for brevity. Only the top two funding segments are quantitatively important: they explain 75% and 10% of the total exposure, respectively, while the third contributes only 3%.

Inspecting the weights, Segment 1 loads negatively on all non-USD currencies—with especially large magnitudes for EUR, GBP, and JPY—consistent with USD funding vis-à-vis all other currencies. For interpretation, we normalize the USD weight to +1, so a $1 position in Segment 1 corresponds to $1 of Dollar funding demand. Segment 2 features a large positive weight on Pan-European currencies (EUR and GBP) and a large negative weights on other major currencies (JPY and USD), consistent with EUR funding vis-à-vis all other currencies. Thus, we normalize the EUR weight to +1, so a $1 position in Segment 2 corresponds to $1 of EUR funding demand. These normalizations fix units and direction only; orthogonality and explanatory power are unchanged because weights are identified up to scale and sign.

We also note that in contrast to the first few price PCs, neither of the two most important independent funding segments features the "carry" currencies (AUD, NZD). This is because there are limited direct funding exposures to these currencies. The logic is straightforward: the appeal of carry is the unhedged exposure, whereas FX funding hedges the FX leg and does not confer the high local-currency yield. For example, funding AUD neither takes FX risk nor earns the AUD carry, so the trade is simply less attractive. Put differently, that direct AUD funding is more expensive than synthetic AUD funding (a positive CIP deviation) likely reflects the limited appeal of the synthetic trade itself. Looking at prices alone might suggest this market is important; combining prices with direct exposures shows

Figure 6: Pricing of funding segments at 1-month tenor



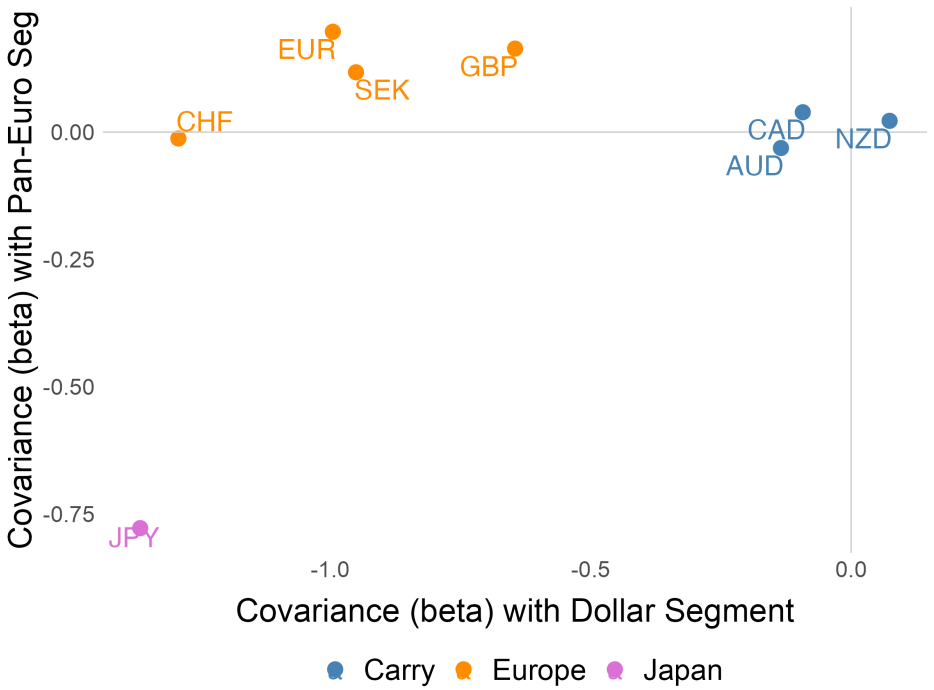
Notes: Prices of funding segments are computed as $P^{\text{seg}}_{k,t} = \sum_n w_{n,k}\, P_{n,t}$ using the weights in Table 2. Series are shown as 30-day moving averages and reported in basis points.

it is not.

Moreover, Appendix Table A1 reports a standard PCA on the direct exposure data. We find that the leading components align with the currency pairs that carry the largest outstanding positions (for example, EUR/USD, GBP/USD, and JPY/USD), which explain most of the variation in direct exposures. In other words, the PCA mainly summarizes concentration in large pairs rather than isolating distinct economic exposures to time-varying funding conditions, since it does not incorporate price changes.

Figure 6 plots the 1M-tenor price of the two funding segments $P^{\text{seg}}_{1,t}$ and $P^{\text{seg}}_{2,t}$. Each series is a portfolio-weighted aggregation of currency-level prices using the weights $w_{n,k}$ from Table 2. By construction, the changes in the two series are orthogonal, so the figure shows distinct dynamics: Segment 1 displays recurrent spikes around stress and quarter-ends, whereas Segment 2 alternates between positive and negative phases—sometimes moving in the opposite direction to Segment 1 (e.g., at quarter-ends) and sometimes in the same direction during global stress episodes (e.g., early 2020). Both series spend most of the sample above zero, indicating that, on average, obtaining FX funding for USD and EUR is more expensive than direct funding in the base currency. The separation underscores that tightening of funding conditions in one segment need not coincide with tightening in the other.

Figure 7: Currency exposures to independent funding segments



Notes: Currency exposures $(\beta_{n,1}, \beta_{n,2})$ come from projecting each currency's price change onto the two funding segments at 1M tenor. The x-axis is exposure to Dollar segment; the y-axis is exposure to Pan-European segment.

6.2 Decomposing CIP deviations

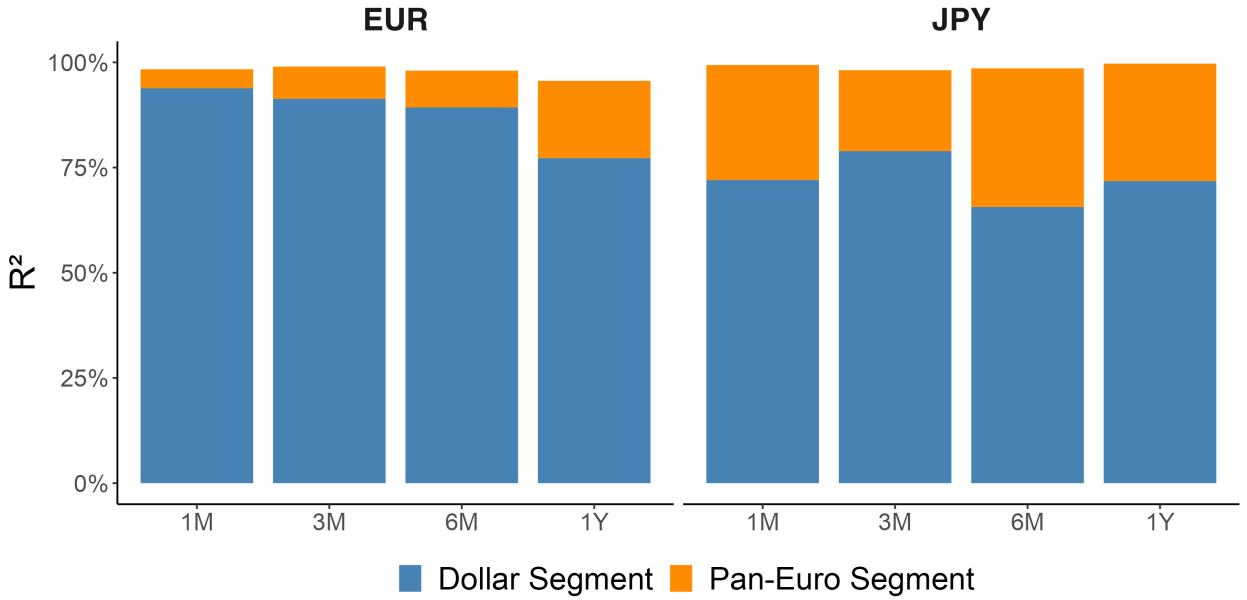
In this subsection, we use the independent funding segments to understand the cross-currency drivers of CIP deviations.

Specifically, Figure 7 plots the estimated coefficients $(\beta_{n,1}, \beta_{n,2})$ from (**??**) at 1M tenor, which measure how each currency's CIP price changes load on the two segments. Three groupings emerge. Pan-European currencies (CHF, EUR, GBP, SEK) lie close to $(-1, 0)$: they load roughly one-for-one on the Dollar segment and only weakly on the Pan-European segment. For example, EUR sits near $(-1, 0.25)$, so demanding \$1 of EUR funding against USD is equivalent to supplying \$1 in the Dollar segment and demanding about \$0.25 in the Pan-European segment in terms of exposure to time-varying funding conditions. JPY sits near $(-1, -1)$, indicating heavy exposure to both segments. By contrast, the carry currencies (AUD, CAD, NZD) cluster around $(0, 0)$, with little exposure to either segment.[4]

These results show that not all dollar funding is the same. For example, borrowing USD against AUD and borrowing USD against EUR both increase USD direct exposure. Yet

[4]The estimated β coefficients are stable over time. Appendix Figure **??** plots five-year rolling-window estimates that show only modest variation across windows.

Figure 8: Segment contributions to currency-specific CIP deviations by tenor



Notes: For each currency and tenor, the bars report the share of explained variation (R^2) in that currency's CIP deviation that is attributable to the Dollar and Pan-European segments, based on the exposure regression in (**??**). Results are shown for EUR and JPY at the 1M, 3M, 6M, and 1Y tenors.

the associated economic exposures differ sharply: the USD/AUD trade has little exposure to either funding segment, whereas the USD/EUR trade is heavily exposed to the Dollar segment and, to a lesser extent, to the Pan-European segment.

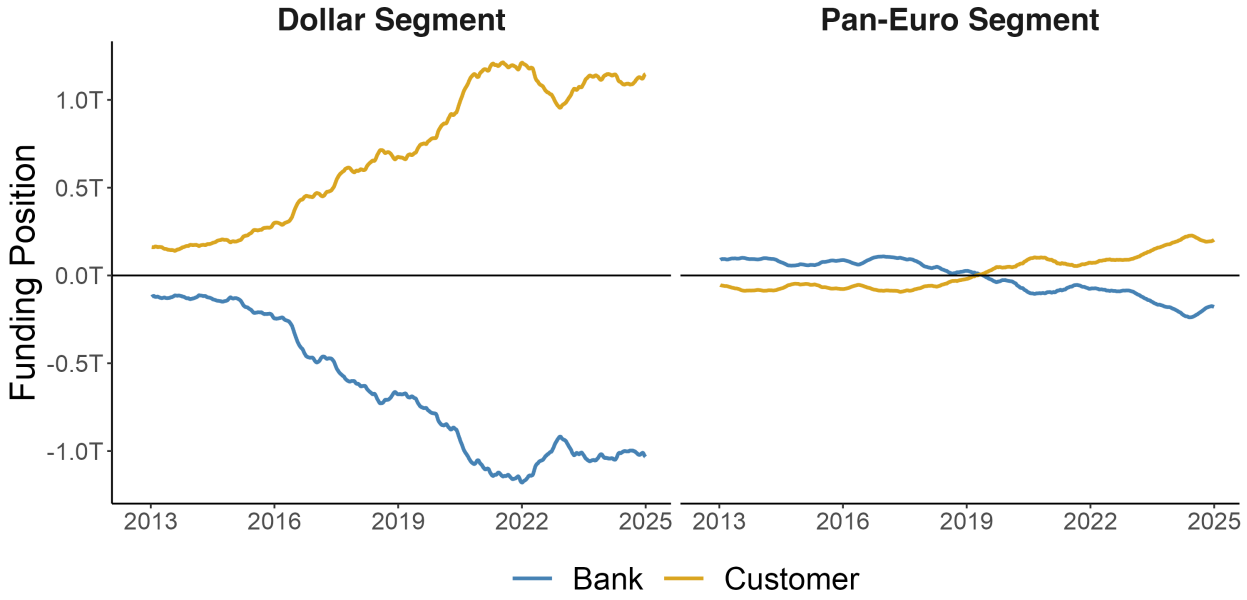
Figure 8 shows that even for a given currency, the relative contribution of the two segments varies systematically with the funding horizon. For EUR, movements in the USD/EUR basis at the 1-month horizon are about 94% attributable to the Dollar segment and 5% to the Pan-European segment; at the 1-year horizon the split shifts to roughly 77% versus 18%. For JPY, both segments contribute materially at all horizons, consistent with its strong loading on each segment. Together, these patterns underscore that the economic forces driving CIP deviations are horizon dependent and segment specific.

6.3 The Economic Exposure to each Funding Segment

This subsection documents investors' economic exposures to each funding segment.

We aggregate economic exposures $\mathcal{E}^{\text{seg}}_{i,k,t,\tau}$ across all banks and across all customers to compare who demands and who supplies funding in the two segments at the 1-month tenor. Figure 9 shows the resulting 90-day-smoothed series. In Segment 1 (Dollar funding), customers are persistent net demanders and banks are the mirror-image suppliers: customer

Figure 9: Economic exposures: banks vs. customers



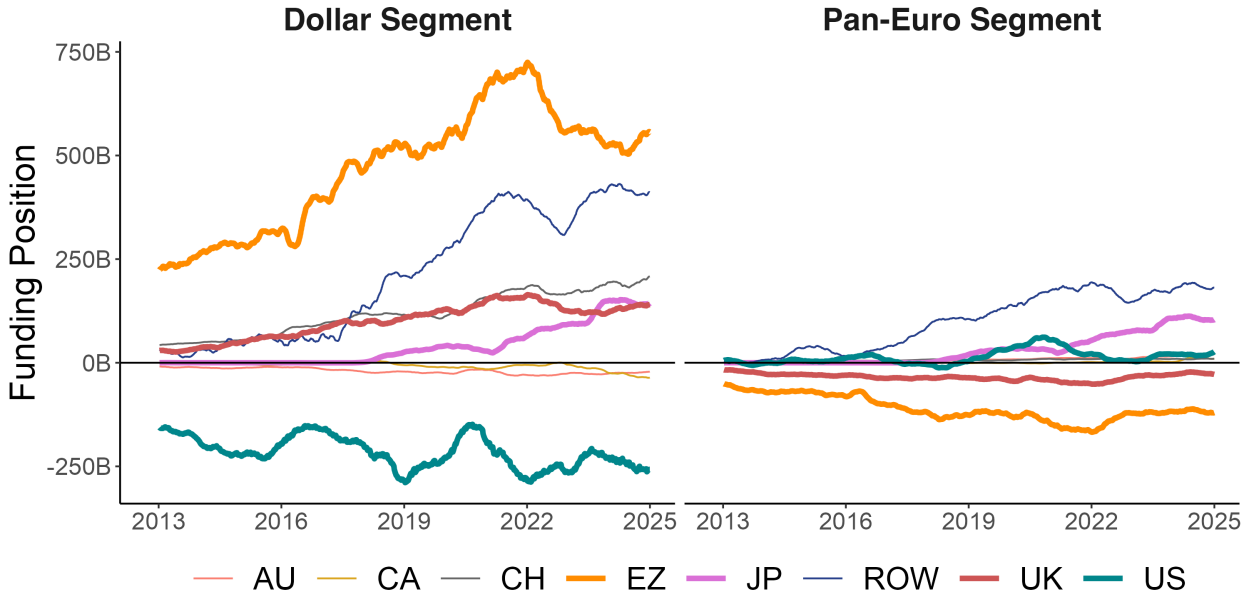
Notes: For each segment $k \in \{1, 2\}$ (Dollar, Pan-European) and $\tau = 1$M, we sum $\mathcal{E}^{\text{seg}}_{i,k,t,\tau}$ across all banks and across all customers. Positive values indicate net demand for the segment; negative values indicate net supply. Series are 90-day moving averages.

demand rises from near zero in 2013 to a little over $1 trillion, while bank supply deepens to roughly -$1 trillion by 2022. In Segment 2 (Pan-European funding), positions are much smaller and switch sign over time: banks move from mild net demand pre-2018 to net supply after 2019, whereas customers transition from mild net supply to net demand; by 2024, customer demand is around $0.2 trillion and bank supply of similar magnitude. The picture is consistent with banks intermediating dollar funding to customers, while the Pan-European segment is closer to balanced and exhibits role reversals over time.

Figure 10 breaks out customers' economic exposures by geographic location and shows stark, systematic patterns across the two segments at 1M tenor. In Segment 1, euro-area (EZ) customers are the dominant net demanders: their dollar funding need rises from roughly $250B in 2013 to a peak near $700B around 2020, and then settles in the $500B range by 2024. U.K., Swiss, and Japanese (UK, CH, JP) customers also post steadily growing dollar demand, each reaching on the order of $150B by the end of the sample. Customers from the Rest of World (ROW) become increasingly important demanders over time, climbing from near zero in 2013 to roughly $300B by 2024. By contrast, U.S. customers are consistently on the supply side, with a persistent negative position that is about $200B.

In Segment 2, the picture flips for several groups. EZ customers are net suppliers of Pan-European funding throughout, with positions generally about -$100B. ROW customers

Figure 10: Economic exposures: Customers



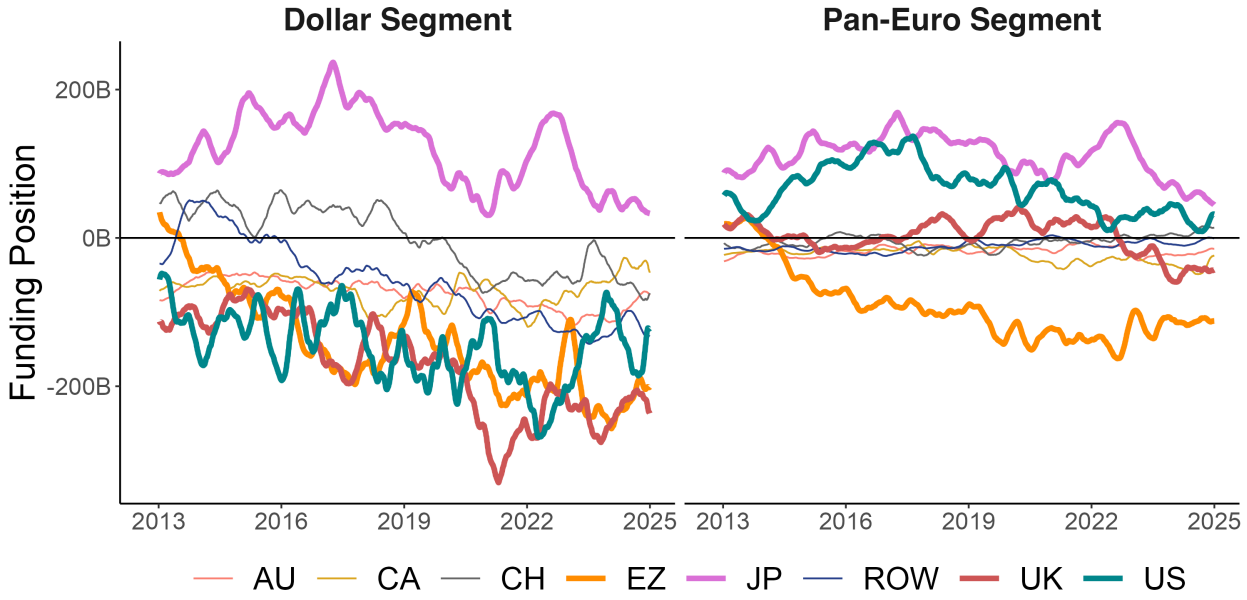
Notes: For each segment $k \in \{1, 2\}$ (Dollar, Pan-European) and $\tau = 1\text{M}$, we show $\mathcal{E}^{\text{seg}}_{i,k,t,\tau}$ for all customers. Positive values indicate net demand for the segment; negative values indicate net supply. Series are 90-day moving averages.

are the largest EUR demanders, rising steadily to around $200B by 2024. JP customers also demand EUR, trending up to roughly $100B in recent years. US customers are small net demanders on average—hovering around zero for much of the sample and turning modestly positive after 2020. Taken together, the panel highlights a clear division of roles: EZ customers demand dollars in Segment 1 and supply euros in Segment 2; ROW and JP customers increasingly demand funding on both fronts; and US customers most often supply dollars while demanding little euro overall.

Figure 11 breaks down the 1M-tenor economic exposures of CLS-settling banks by geographic location (non-settling banks are small and omitted). In Segment 1, JP banks are persistent net demanders of dollars, typically in the $100-$200B range, peaking a little above $200B in 2016-2018 and again around 2020-2021. By contrast, UK, US, and EZ banks are the main net suppliers: each around $100-200B on average. Other regions are small by comparison. In Segment 2, the signs flip. US and JP banks are net demanders of euro funding, on the order of $50 to $150B. EZ and UK banks are net suppliers: EZ banks provide roughly $100B, while UK banks move from near zero before 2021 to about $50B supply thereafter. Other regions are small by comparison.

Taken together, Dollar segment funding is largely intermediated by UK/US/EZ banks, whereas Pan-European Segment funding is sourced primarily in the euro area and supple-

Figure 11: Economic exposures: CLS-settling banks



Notes: For each segment $k \in \{1, 2\}$ (Dollar, Pan-European) and $\tau = 1\text{M}$, we show $\mathcal{E}^{\text{seg}}_{i,k,t,\tau}$ for all CLS-settling banks by geographic location. Positive values indicate net demand for the segment; negative values indicate net supply. Series are 90-day moving averages.

mented by UK supply to meet demand from US and Japanese banks. These patterns also show that large global banks diversify across segments. For example, US banks supply funding in Segment 1 while demanding funding in Segment 2. Although both actions raise their direct exposure to dollar supply, this allocation spreads their economic exposure across segments whose price changes are orthogonal by construction.

6.4 Relative Dominance of Dollar vs. Pan-European Funding Across Horizons and Over Time

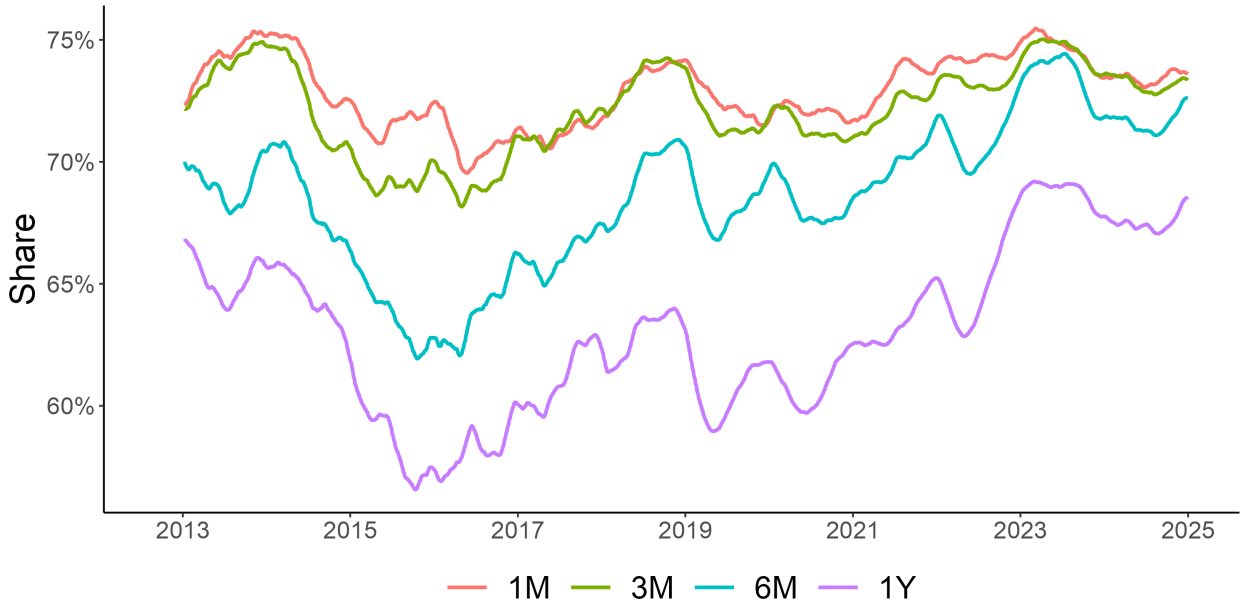
To compare the importance of Dollar vs. Pan-European funding across horizons and over time, we measure total economic exposures for each segment. Let

$$\sigma_k = \text{sd}(\Delta P^{\text{seg}}_{k,t}) \tag{30}$$

be the standard deviation of price changes in segment $k \in \{1, 2\}$. For date t and horizon τ, define the total economic exposures across all investors i of segment k as

$$A_{k,t,\tau} = \sigma_k \sum_i |\mathcal{E}^{\text{seg}}_{i,k,t,\tau}|. \tag{31}$$

Figure 12: Relative dominance of Dollar vs. Pan-European funding



Notes: The series plots $\text{DollarImportance}_{t,\tau}$ defined in (32) for each tenor τ. Curves are 90-day moving averages.

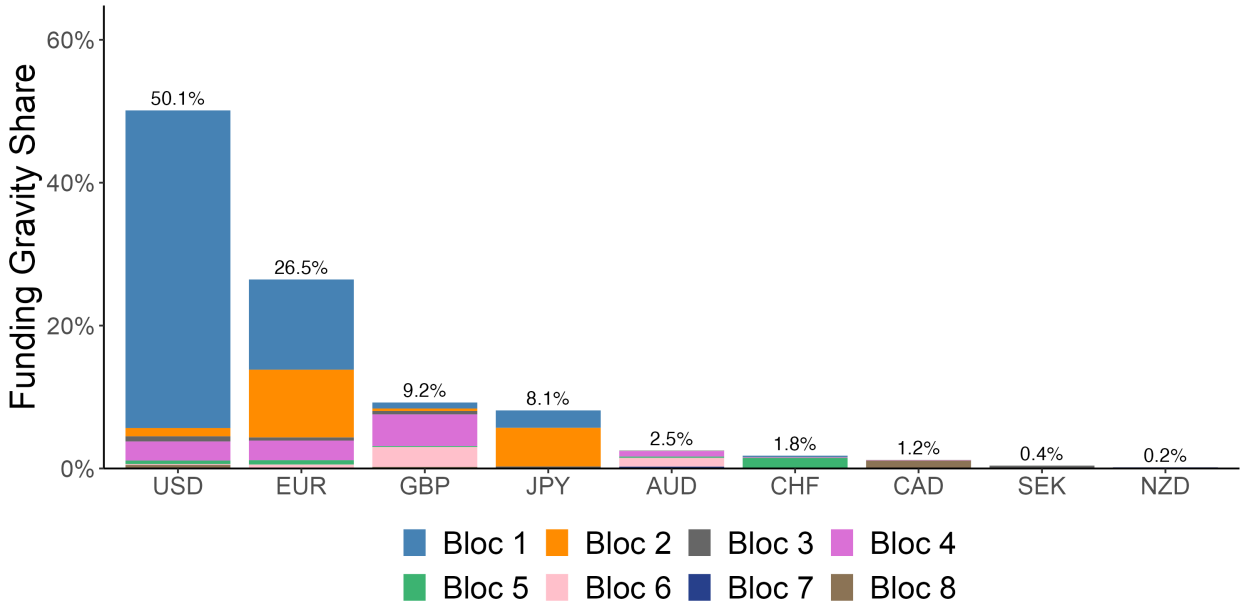
We take absolute values because market clearing implies that investor exposures sum to zero; $|\cdot|$ aggregates the gross amount of exposure. Multiplying by σ_k makes $A_{k,t,\tau}$ invariant to the arbitrary scaling of segment weights (since $|\mathcal{E}^{\text{seg}}_{i,k,t,\tau}|$ scales inversely while $\Delta P^{\text{seg}}_{k,t}$ scales proportionally). We then measure the relative dominance of the Dollar funding (Segment 1) vs. the Pan-European funding (Segment 2) as

$$\text{DollarImportance}_{t,\tau} = \frac{A_{1,t,\tau}}{A_{1,t,\tau} + A_{2,t,\tau}}. \tag{32}$$

By construction, $\text{DollarImportance}_{t,\tau} \in [0,1]$; the complementary share $1-\text{DollarImportance}_{t,\tau}$ is the EUR-funding dominance.

Figure 12 plots $\text{DollarImportance}_{t,\tau}$ for $\tau \in \{1M, 3M, 6M, 1Y\}$. To reduce high-frequency noise, we display 90-day moving averages. Comparing total economic exposures across all investors, the Dollar segment averages about 70% of combined activity. However, this dominance declines with funding horizons: roughly 75% at 1M and 3M, falling to around 65% at 6M and 1Y. There is also notable time-series variation, whereby dollar dominance dipped around 2016, possibly related to the U.S. money-market fund reform that tightened access to dollar funding.

Figure 13: Country's Funding Gravity and Its Sources



Notes: This figure plots each country's funding gravity, with total labeled on top of the corresponding bar. The segments show the contribution from the eight orthogonal funding blocs.

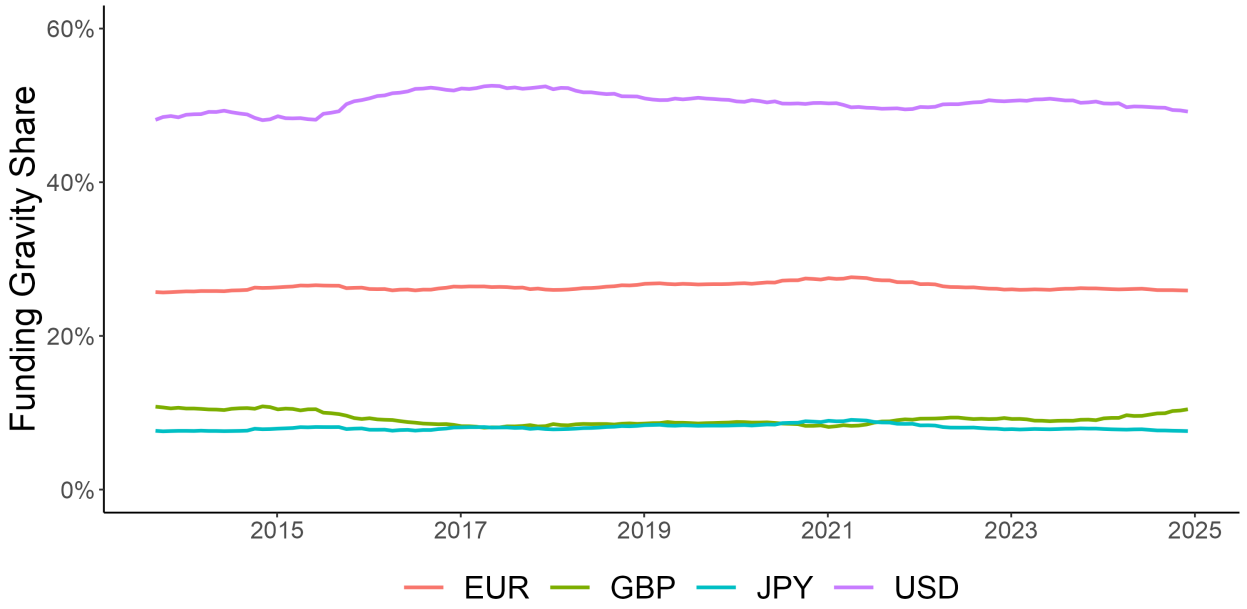
7 Funding Gravity of Countries

In this section, we quantify each country's funding gravity, defined as its capacity to attract global capital when its funding conditions improve. We measure funding gravity by aggregating across the independent funding blocs identified earlier, weighted by each country's participation in those blocs.

Figure 13 shows that funding gravity is highly concentrated across currencies. The USD accounts for 50.1% of total gravity, roughly equal to all other currencies combined. The EUR follows with 26.5%, while the GBP and the JPY contribute 9.2% and 8.1%, respectively. The remaining currencies together represent less than 6% of the total. The stacked bars decompose each country's funding gravity into contributions from the eight orthogonal funding blocs. Most of the USD and EUR shares are explained by the two dominant blocs, Dollar and Pan-European, while the remaining blocs contribute marginally. This pattern highlights that global FX funding is organized primarily around two independent ecosystems that anchor international financial power.

Figure 14 plots the evolution of funding gravity in our sample period. The hierarchy observed in the cross-section remains remarkably stable over time. The relative ordering among the major currencies has shown little change over more than a decade, underscoring the persistence of global funding structures and the enduring dominance of the Dollar and

Figure 14: Time-series of Funding Gravity



Notes: This figure plots the time-series funding gravity for USD, EUR, GBP, and JPY.

Pan-European blocs in shaping international capital flows.

8 Conclusion

We develop a new framework to quantify countries' financial power. Our measure, funding gravity, captures a country's structural influence in global finance through its ability to attract capital when its own funding conditions change. Applying this framework to the global FX funding market, where we leverage a proprietary data of eighty agent groups in eight geographies, we identify distinct funding blocs that generate orthogonal shocks to funding conditions and uncover a persistent hierarchy of financial power. The United States dominates this hierarchy, followed by the euro area, the United Kingdom, and Japan. The stability of this structure over more than a decade reflects the enduring capital pull of the major currencies and the concentration of global funding activity in a few key blocs.

Beyond the FX market, the concept of funding gravity offers a general approach to studying financial dominance and capital flow interdependence across markets. It can be used to examine how financial power shifts in response to policy, crisis, or technological change, and how it shapes global risk transmission. More broadly, our results highlight that financial power is not merely exercised, but it is embedded in the structure of global funding itself.

References

Abadi, J., J. Fernàndez-Villaverde, and D. Sanches. 2025. International currency dominance. Working Paper.

An, Y. 2025. An arbitrage foundation for demand effects in asset pricing. Working paper, Johns Hopkins University.

An, Y., and A. W. Huber. 2025. Demand propagation through traded risk factors. Working Paper.

Bianchi, J., S. Horn, G. Rosso, and C. Sosa-Padilla. 2025. Financial cooperation in a fragmented world. Presentation slides, presented at the IMF–Kiel Institute Conference on Geoeconomic Fragmentation, October 2025.

Broner, F., A. Martin, J. Meyer, and C. Trebesch. 2025. Hegemonic globalization. Working paper.

Cenedese, G., P. Della Corte, and T. Wang. 2021. Currency mispricing and dealer balance sheets. *The Journal of Finance* 76:2763–803.

Clayton, C., A. Coppola, M. Maggiori, and J. Schreger. 2025. Geoeconomic pressure. Working paper.

Clayton, C., M. Maggiori, and J. Schreger. 2025. A framework for geoeconomics. *Econometrica* Forthcoming.

Coppola, A., A. Krishnamurthy, and C. Xu. 2025. Liquidity, debt denomination, and currency dominance. Working Paper.

Dao, M. C., P.-O. Gourinchas, and O. Itskhoki. 2025. Breaking parity: Equilibrium exchange rates and currency premia. Working Paper.

De Leo, P., L. Keller, and D. Zou. 2025. Speculation, forward exchange demand, and cip deviations in emerging economies. Working Paper.

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

Du, W., and A. W. Huber. 2025. USD asset holding and hedging around the globe. Working Paper.

Du, W., G. Strasser, and A. Verdelhan. 2025. Repo and fx swap: A tale of two markets. Working Paper.

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

European Central Bank. 2025. The international role of the euro. https://www.ecb.europa.eu/press/other-publications/ire/html/index.en.html.

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

Kloks, P., E. Mattille, and A. Ranaldo. 2024. Hunting for dollars. Working Paper.

Matsuyama, K., N. Kiyotaki, and A. Matsui. 1993. Toward a theory of international currency. *The Review of Economic Studies* 60:283–307.

Miranda-Agrippino, S., and H. Rey. 2020. U.S. Monetary Policy and the Global Financial Cycle. *The Review of Economic Studies* 87:2754–76.

Moskowitz, T. J., C. P. Ross, S. Y. Ross, and K. Vasudevan. 2024. Quantities and covered-

interest parity. Working Paper.

Shin, H. S., P. Wooldridge, and D. Xia. 2025. US dollar's slide in April 2025: the role of FX hedging. Working Paper, Bank for International Settlements.

Zhang, C. 2014. An information-based theory of international currency. *Journal of International Economics* 93:286–301.

A Proof

This appendix provide proofs omitted in the main text.

A.1 Proof of Proposition 1

Let $\tilde{\mathbf{W}} \in \mathbb{R}^{(N+1)\times N}$ collect the weights $w_{n,k}$, stack variables as $\mathbf{a}_t, \mathbf{f}_t, \tilde{\mathbf{f}}_t, \mathbf{Q}_t, \mathbf{P}_t^{\text{seg}} \in \mathbb{R}^{N+1}$, and let

$$\boldsymbol{\Gamma} := \theta \text{diag}(\mu_1, \ldots, \mu_N). \tag{A1}$$

By definition of segment prices,

$$\mathbf{P}_t^{\text{seg}} = -\tilde{\mathbf{W}}^\top \mathbf{f}_t. \tag{A2}$$

Using the demand specification, aggregate demand stacks as

$$\mathbf{Q}_t = \mathbf{a}_t - \tilde{\mathbf{W}}\boldsymbol{\Gamma}\mathbf{P}_t^{\text{seg}} = \mathbf{a}_t + \tilde{\mathbf{W}}\boldsymbol{\Gamma}\tilde{\mathbf{W}}^\top \mathbf{f}_t = \mathbf{a}_t + \mathbf{L}\mathbf{f}_t, \tag{A3}$$

where

$$\mathbf{L} := \tilde{\mathbf{W}}\boldsymbol{\Gamma}\tilde{\mathbf{W}}^\top, \qquad (\mathbf{L})_{nn} = \theta \sum_{k=1}^{N} \mu_k w_{n,k}^2 = \theta\varphi_n. \tag{A4}$$

Country n chooses $f_{n,t}$ to maximize

$$Q_{n,t} - \frac{c}{2}\big(f_{n,t} - \tilde{f}_{n,t}\big)^2, \tag{A5}$$

taking $\{f_{m,t}\}_{m\neq n}$ as given. From (A3) and (A4),

$$\frac{\partial Q_{n,t}}{\partial f_{n,t}} = (\mathbf{L})_{nn} = \theta\varphi_n. \tag{A6}$$

The first-order condition is therefore

$$0 = \theta\varphi_n - c\big(f_{n,t} - \tilde{f}_{n,t}\big), \tag{A7}$$

which solves to (7). Uniqueness follows since, for each n, the objective in (A5) is strictly concave in $f_{n,t}$ with second derivative $-c < 0$, and the best response in (7) is a singleton independent of rivals' choices.

A.2 Proof of Proposition 2

Let $\Delta\mathbf{P}_{0,t} \in \mathbb{R}^N$ collect bilateral price changes against the numeraire,

$$\Delta\mathbf{P}_{0,t} := \big(\Delta P_{1,0,t}, \ldots, \Delta P_{N,0,t}\big)^\top, \tag{A8}$$

and let $\Delta\mathbf{P}_t^{\text{seg}} \in \mathbb{R}^N$ and $\Delta\mathbf{Q}_t^{\text{seg}} \in \mathbb{R}^N$ stack segment price and demand changes, respectively. Denote by $\mathbf{W} \in \mathbb{R}^{N\times N}$ the weight matrix with entries $w_{n,k}$ for $n, k = 1, \ldots, N$. From (9) and the segment aggregation $Q_{n,t} = \sum_{k=1}^N w_{n,k} Q_{k,t}^{\text{seg}}$,

$$\Delta\mathbf{P}_t^{\text{seg}} = \mathbf{W}^\top \Delta\mathbf{P}_{0,t}, \qquad \Delta\mathbf{Q}_t = \mathbf{W}\Delta\mathbf{Q}_t^{\text{seg}}, \tag{A9}$$

where $\Delta\mathbf{Q}_t \in \mathbb{R}^N$ stacks $\Delta Q_{n,t}$ for $n = 1, \ldots, N$.

Assume that $\operatorname{var}(\Delta\mathbf{P}_{0,t})$ is nonsingular. (The more general case where $\operatorname{var}(\Delta\mathbf{P}_{0,t})$ is singular is analyzed in An and Huber (2025).) Take its Cholesky factorization

$$\operatorname{var}(\Delta\mathbf{P}_{0,t}) = \mathbf{U}^\top\mathbf{U}, \tag{A10}$$

with $\mathbf{U}$ upper triangular and $\operatorname{diag}(\mathbf{U}) > 0$. Consider the symmetric matrix $\mathbf{U}\operatorname{var}(\Delta\mathbf{Q}_t)\mathbf{U}^\top$ and take its spectral decomposition

$$\mathbf{U}\operatorname{var}(\Delta\mathbf{Q}_t)\mathbf{U}^\top\mathbf{G} = \mathbf{G}\mathbf{\Pi}, \tag{A11}$$

where $\mathbf{G}$ is orthonormal ($\mathbf{G}^\top\mathbf{G} = \mathbf{I}_N$) and $\mathbf{\Pi} = \operatorname{diag}(\pi_1, \ldots, \pi_N)$ has nonnegative entries. Define

$$\mathbf{W} := \mathbf{U}^{-1}\mathbf{G}. \tag{A12}$$

Then, using (A9) and (A10),

$$\operatorname{var}(\Delta\mathbf{P}_t^{\text{seg}}) = \mathbf{W}^\top \operatorname{var}(\Delta\mathbf{P}_{0,t})\mathbf{W} = \mathbf{G}^\top(\mathbf{U}^\top)^{-1}\mathbf{U}^\top\mathbf{U}\mathbf{U}^{-1}\mathbf{G} = \mathbf{I}_N. \tag{A13}$$

Hence $\Delta P_{k,t}^{\text{seg}}$ are mutually uncorrelated (and normalized) across k, which establishes (12).

Next, from (A9),

$$\operatorname{var}(\Delta\mathbf{Q}_t) = \mathbf{W}\operatorname{var}(\Delta\mathbf{Q}_t^{\text{seg}})\mathbf{W}^\top. \tag{A14}$$

Left- and right-multiply (A14) by $\mathbf{U}$ and $\mathbf{U}^\top$, and use (A12):

$$\mathbf{U}\operatorname{var}(\Delta\mathbf{Q}_t)\mathbf{U}^\top = (\mathbf{U}\mathbf{W})\operatorname{var}(\Delta\mathbf{Q}_t^{\text{seg}})(\mathbf{U}\mathbf{W})^\top = \mathbf{G}\operatorname{var}(\Delta\mathbf{Q}_t^{\text{seg}})\mathbf{G}^\top. \tag{A15}$$

Comparing (A15) with (A11) and using orthonormality of $\mathbf{G}$,

$$\text{var}(\Delta \mathbf{Q}_t^{\text{seg}}) = \mathbf{\Pi}. \tag{A16}$$

Therefore $\Delta Q_{k,t}^{\text{seg}}$ are mutually uncorrelated across k, which establishes (13).

B Generalization to Inelastic Market and Investor Heterogeneity

The baseline model assumes a perfectly elastic market, meaning that funding prices do not adjust with the level of aggregate demand. Investor heterogeneity is also abstracted away, represented by a single observable net demand per country from all other countries. This appendix extends the framework to incorporate both market inelasticity and heterogeneous investors. We show that, to obtain uncorrelated funding shocks and uncorrelated attractiveness shocks as in Proposition 2, it is equivalent to jointly orthogonalize equilibrium price changes and quantity changes aggregated across investors.

We consider I investors in total. The equilibrium quantity change (including both the demand and supply of funding) of investor i in country n at time t is denoted by $\Delta Q_{i,n,t}$. Following the notation in the main text, for any given weight matrix $\mathbf{W} := \{w_{n,k}\}_{n,k=1,\dots,N}$ and its inverse $\mathbf{W}^{-1}$, whose (k,n)-th element is $v_{k,n}$, the equilibrium quantity change of investor i in segment k at time t is defined as

$$\Delta Q_{i,k,t}^{\text{seg}} = \sum_{n=1}^{N} v_{k,n} \Delta Q_{i,n,t}. \tag{A17}$$

Generalizing equation (11), we define the change in the demand of investor i for segment k in period t as

$$\Delta \tilde{Q}_{i,k,t}^{\text{seg}} = \gamma_{i,k} \sum_{n=1}^{N} v_{k,n} \Delta a_{n,t} - \theta \mu_{i,k} \Delta P_{k,t}^{\text{seg}}. \tag{A18}$$

The key difference relative to the baseline equation (11) is that each investor i can have a distinct loading $\gamma_{i,k}$ on the segment-level attractiveness component $\sum_{n=1}^{N} v_{k,n} \Delta a_{n,t}$, allowing for heterogeneity across both investors and segments. Similarly, the amount of capital $\mu_{i,k}$ investor i allocates to segment k can also vary across investors and segments.

We assume that the equilibrium quantity change combines both the demand and supply of funding:

$$\Delta Q_{i,k,t}^{\text{seg}} = \Delta \tilde{Q}_{i,k,t}^{\text{seg}} - \zeta_{i,k} \big(\Delta P_{k,t}^{\text{seg}} - \Delta \tilde{P}_{k,t}^{\text{seg}} \big), \tag{A19}$$

where the change in the underlying funding condition of segment k is defined as in equa-

tion (9),

$$\Delta \tilde{P}^{\text{seg}}_{k,t} = \sum_{n=1}^{N} w_{n,k}\big(\Delta \tilde{f}_{0,t} - \Delta \tilde{f}_{n,t}\big). \tag{A20}$$

The parameter $\zeta_{i,k}$ captures the supply elasticity of investor i with respect to segment k. A higher $\zeta_{i,k}$ indicates that investor i is more willing to supply funding in response to a price dislocation between the observed equilibrium price change $\Delta P^{\text{seg}}_{k,t}$ and the component $\Delta \tilde{P}^{\text{seg}}_{k,t}$ driven by underlying funding condition changes. The baseline model of a perfectly elastic market corresponds to the limiting case $\zeta_{i,k} \to \infty$.

Market clearing for segment k across all investors implies

$$\sum_{i=1}^{I} \Delta Q^{\text{seg}}_{i,k,t} = 0. \tag{A21}$$

Summing equation (A19) over all investors then yields

$$\Delta P^{\text{seg}}_{k,t} = \Delta \tilde{P}^{\text{seg}}_{k,t} + \frac{\sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,k,t}}{\sum_{i=1}^{I} \zeta_{i,k}}. \tag{A22}$$

Equation (A22) shows that the price change of segment k equals the underlying funding condition shock $\Delta \tilde{P}^{\text{seg}}_{k,t}$ plus an additional term reflecting the aggregate demand shock $\sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,k,t}$, scaled by the aggregate elasticity $\sum_{i=1}^{I} \zeta_{i,k}$. A larger value of $\sum_{i=1}^{I} \zeta_{i,k}$ implies that the market is more elastic and therefore dampens the price impact of a given demand shock. Plugging (A18) into (A22), one can explicitly express the equilibrium price change as the underlying shocks as

$$\Delta P^{\text{seg}}_{k,t} = \frac{\sum_{i=1}^{I} \zeta_{i,k}}{\sum_{i=1}^{I} \zeta_{i,k} + \theta \sum_{i=1}^{I} \mu_{i,k}} \Delta \tilde{P}^{\text{seg}}_{k,t} + \frac{\sum_{i=1}^{I} \gamma_{i,k}}{\sum_{i=1}^{I} \zeta_{i,k} + \theta \sum_{i=1}^{I} \mu_{i,k}} \sum_{n=1}^{N} v_{k,n} \Delta a_{n,t}. \tag{A23}$$

Plugging (A23) into (A18), the aggregate demand shock for segment k is

$$\sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,k,t} = \frac{\sum_{i=1}^{I} \zeta_{i,k}}{\sum_{i=1}^{I} \zeta_{i,k} + \theta \sum_{i=1}^{I} \mu_{i,k}} \left(\sum_{i=1}^{I} \gamma_{i,k} \sum_{n=1}^{N} v_{k,n} \Delta a_{n,t} - \theta \sum_{i=1}^{I} \mu_{i,k} \Delta \tilde{P}^{\text{seg}}_{k,t} \right) \tag{A24}$$

Proposition 2 shows that one can choose the weights $w_{n,k}$ to jointly orthogonalize the underlying funding condition shocks $\Delta \tilde{P}^{\text{seg}}_{k,t}$ and the aggregate demand shocks $\sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,k,t}$. However, these objects are not directly observable. We next show that these orthogonality

requirements are equivalent to

$$\operatorname{cov}\left(\Delta P^{\text{seg}}_{k,t}, \Delta P^{\text{seg}}_{l,t}\right) = 0 \quad \text{for any } k \neq j, \tag{A25}$$

$$\sum_{i=1}^{I} \operatorname{cov}\left(\Delta Q^{\text{seg}}_{i,k,t}, \Delta Q^{\text{seg}}_{i,l,t}\right) = 0 \quad \text{for any } k \neq j. \tag{A26}$$

In other words, the same weights $w_{n,k}$ can be chosen to jointly orthogonalize equilibrium price and quantity changes aggregated across investors. The covariance of quantity changes in equation (A26) is computed within each investor: for a given investor i, we first compute $\operatorname{cov}(\Delta Q^{\text{seg}}_{i,k,t}, \Delta Q^{\text{seg}}_{i,l,t})$ over time t, and then aggregate across investors (by summing, or equivalently, by averaging up to a constant).

First, we consider the weights $\{w_{n,k}\}$ such that

$$\operatorname{cov}(\Delta \tilde{P}^{\text{seg}}_{k,t}, \Delta \tilde{P}^{\text{seg}}_{l,t}) = 0 \quad \text{for any } k \neq j, \tag{A27}$$

$$\operatorname{cov}\left(\sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,k,t}, \sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,l,t}\right) = 0 \quad \text{for any } k \neq j. \tag{A28}$$

Assumption (14), together with equations (A20), (A24), (A27), and (A28) implies

$$\operatorname{cov}\left(\sum_{n=1}^{N} w_{n,k}\left(\Delta \tilde{f}_{0,t} - \Delta \tilde{f}_{n,t}\right), \sum_{n=1}^{N} w_{n,l}\left(\Delta \tilde{f}_{0,t} - \Delta \tilde{f}_{n,t}\right)\right) = 0 \quad \text{for any } k \neq j, \tag{A29}$$

$$\operatorname{cov}\left(\sum_{n=1}^{N} v_{k,n} \Delta a_{n,t}, \sum_{n=1}^{N} v_{l,n} \Delta a_{n,t}\right) = 0 \quad \text{for any } k \neq j, \tag{A30}$$

$$\operatorname{cov}\left(\Delta \tilde{P}^{\text{seg}}_{k,t}, \sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,l,t}\right) = \operatorname{cov}\left(\Delta \tilde{P}^{\text{seg}}_{l,t}, \sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,k,t}\right) = 0 \quad \text{for any } k \neq j. \tag{A31}$$

Hence, (A22) implies

$$\begin{aligned}\operatorname{cov}(\Delta P^{\text{seg}}_{k,t}, \Delta P^{\text{seg}}_{l,t}) &= \operatorname{cov}(\Delta \tilde{P}^{\text{seg}}_{k,t}, \Delta \tilde{P}^{\text{seg}}_{l,t}) + \operatorname{cov}\left(\frac{\sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,k,t}}{\sum_{i=1}^{I} \zeta_{i,k}}, \frac{\sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,l,t}}{\sum_{i=1}^{I} \zeta_{i,l}}\right) \\ &+ \operatorname{cov}\left(\Delta \tilde{P}^{\text{seg}}_{k,t}, \frac{\sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,l,t}}{\sum_{i=1}^{I} \zeta_{i,l}}\right) + \operatorname{cov}\left(\Delta \tilde{P}^{\text{seg}}_{l,t}, \frac{\sum_{i=1}^{I} \Delta \tilde{Q}^{\text{seg}}_{i,l,t}}{\sum_{i=1}^{I} \zeta_{i,k}}\right) = 0,\end{aligned} \tag{A32}$$

which proves (A25).

Next, observe that for any two investors i and j (whether identical or distinct) and for any pair of different segments $k \neq l$, equations (14), (A18), (A23), (A29), and (A30) together

imply that

$$\operatorname{cov}(\Delta\tilde{Q}^{\text{seg}}_{i,k,t}, \Delta\tilde{Q}^{\text{seg}}_{j,l,t}) = \operatorname{cov}\left(\gamma_{i,k}\sum_{n=1}^{N} v_{k,n}\Delta a_{n,t} - \theta\mu_{i,k}\Delta P^{\text{seg}}_{k,t}, \gamma_{j,l}\sum_{n=1}^{N} v_{l,n}\Delta a_{n,t} - \theta\mu_{j,l}\Delta P^{\text{seg}}_{l,t}\right) = 0. \tag{A33}$$

Using (A19) and (A22), for any investor j we have

$$\Delta Q^{\text{seg}}_{j,k,t} = -\psi_{j,k}\sum_{i=1}^{I}\Delta\tilde{Q}^{\text{seg}}_{i,k,t} + \Delta\tilde{Q}^{\text{seg}}_{j,k,t}, \qquad \psi_{j,k} := \frac{\zeta_{j,k}}{\sum_{i=1}^{I}\zeta_{i,k}}. \tag{A34}$$

Hence,

$$\begin{aligned}\operatorname{cov}\left(\Delta Q^{\text{seg}}_{j,k,t}, \Delta Q^{\text{seg}}_{j,l,t}\right) = &\operatorname{cov}\left(\psi_{j,k}\sum_{i=1}^{I}\Delta\tilde{Q}^{\text{seg}}_{i,k,t}, \psi_{j,l}\sum_{i=1}^{I}\Delta\tilde{Q}^{\text{seg}}_{i,l,t}\right) + \operatorname{cov}(\Delta\tilde{Q}^{\text{seg}}_{j,k,t}, \Delta\tilde{Q}^{\text{seg}}_{j,l,t}) \\ &- \operatorname{cov}\left(\psi_{j,k}\sum_{i=1}^{I}\Delta\tilde{Q}^{\text{seg}}_{i,k,t}, \Delta\tilde{Q}^{\text{seg}}_{j,l,t}\right) - \operatorname{cov}\left(\psi_{j,l}\sum_{i=1}^{I}\Delta\tilde{Q}^{\text{seg}}_{i,l,t}, \Delta\tilde{Q}^{\text{seg}}_{j,k,t}\right).\end{aligned} \tag{A35}$$

Summing across investors gives

$$\begin{aligned}&\sum_{i=1}^{I}\operatorname{cov}(\Delta Q^{\text{seg}}_{i,k,t}, \Delta Q^{\text{seg}}_{i,l,t}) = \left(\sum_{i=1}^{I}\psi_{i,k}\psi_{i,l}\right)\operatorname{cov}\left(\sum_{i=1}^{I}\Delta\tilde{Q}^{\text{seg}}_{i,k,t}, \sum_{i=1}^{I}\Delta\tilde{Q}^{\text{seg}}_{i,l,t}\right) + \sum_{i=1}^{I}\operatorname{cov}(\Delta\tilde{Q}^{\text{seg}}_{i,k,t}, \Delta\tilde{Q}^{\text{seg}}_{i,l,t}) \\ &- \operatorname{cov}\left(\sum_{i=1}^{I}\Delta\tilde{Q}^{\text{seg}}_{i,k,t}, \sum_{i=1}^{I}\psi_{i,k}\Delta\tilde{Q}^{\text{seg}}_{i,l,t}\right) - \operatorname{cov}\left(\sum_{i=1}^{I}\Delta\tilde{Q}^{\text{seg}}_{i,l,t}, \sum_{i=1}^{I}\psi_{i,l}\Delta\tilde{Q}^{\text{seg}}_{i,k,t}\right) = 0,\end{aligned} \tag{A36}$$

which proves (A26).

For the reverse direction of showing that (A25) and (A26) imply (A27) and (A28), one can avoid lengthy algebra by invoking the uniqueness result in An (2025). Specifically, the weights $\{w_{n,k}\}$ identified either from (A25) and (A26) or from (A27) and (A28) exist and are unique up to rotation within the same eigenspace. Since we have already shown that (A27) and (A28) imply (A25) and (A26), uniqueness forces the converse: any weights $\{w_{n,k}\}$ that satisfy (A25) and (A26) must coincide with those recovered from (A27) and (A28).

C Additional Empirical Results

Table A1 runs a PCA on changes in net exposures, $\Delta\mathcal{E}_{i,n,t,\tau} := \mathcal{E}_{i,n,t,\tau} - \mathcal{E}_{i,n,t-1,\tau}$. Fix a tenor τ. For each investor i, compute the time-series covariance of the currency vector $\Delta\mathcal{E}_{i,\cdot,t,\tau}$

Table A1: PCA of changes in net exposure

	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8
AUD	0.00	-0.02	-0.02	-1.00	0.13	0.05	0.01	0.02
CAD	0.01	0.00	-0.02	-0.06	-0.07	-1.00	0.01	0.00
CHF	-0.03	-0.02	-0.02	-0.12	-1.00	0.08	0.02	0.01
EUR	-1.00	0.12	0.02	0.00	0.02	-0.01	0.00	0.00
GBP	-0.12	-1.00	-0.06	0.02	0.02	0.00	0.01	0.00
JPY	-0.01	0.06	-1.00	0.02	0.02	0.02	0.02	0.02
NZD	0.00	0.00	-0.02	-0.03	-0.01	0.00	-0.15	-1.00
SEK	-0.01	0.00	-0.02	-0.01	-0.01	-0.01	-1.00	0.15
USD	1.17	0.85	1.13	1.18	0.92	0.87	1.07	0.79
Var explained	55.4%	21.5%	13.4%	3.9%	3.0%	2.2%	0.4%	0.2%

Notes: PCA is applied to the investor-averaged within-investor covariance matrix $\Sigma_\tau^{\mathcal{E}}$ at $\tau = 1\text{M}$ horizon. "Var explained" is the fraction of the average within-investor covariance of ΔB captured by each PC.

across n, and then average across investors:

$$\Sigma_\tau^{\mathcal{E}} \;=\; \frac{1}{I}\sum_{i=1}^{I} \operatorname{cov}_t\big(\Delta\mathcal{E}_{i,\cdot,t,\tau}\big). \tag{A37}$$

By market clearing, $\sum_n \Delta\mathcal{E}_{i,n,t,\tau} = 0$; we therefore drop USD when estimating $\Sigma_\tau^{\mathcal{E}}$, compute principal components on the remaining currencies, and recover the USD loading in each component as the negative sum of the non-USD loadings.

Table A1 reports the 1M results; other tenors are similar and omitted for brevity. The first three PCs explain $55.4\% + 21.5\% + 13.4\% \approx 90.3\%$ of the average within-investor covariance in changes of net exposures. Each leading PC is essentially one currency against USD: PC1 aligns with EUR/USD, PC2 with GBP/USD, PC3 with JPY/USD; subsequent PCs line up with AUD/USD, CHF/USD, CAD/USD, SEK/USD, and NZD/USD.