Demand Propagation Through Traded Risk Factors*

Yu An[†] Amy W. Huber[‡]

First version: April 2024.
This version: June 2026.

Abstract

We quantify how uninformed demand affects exchange rates in an interconnected FX market. Using 11 years of daily customer-bank FX flows and exchange-rate returns across 17 currencies, we find that a $1 billion demand shock to one currency moves other exchange rates by up to 9 basis points, with substantial heterogeneity across currency pairs. Our key insight is that cross-currency propagation can be decomposed into factor-level repricing and currencies' exposures to common risks. We make this decomposition empirically tractable by identifying three traded risk factors that account for 90% of the non-diversifiable risk banks bear when absorbing customer demand imbalances and that plausibly exhibit no cross-factor price effects. We estimate each factor's price sensitivity using sovereign bond auction announcements as instruments for non-informational shocks to factor demand and find that the FX market is highly elastic. Consistent with the model's predictions, out-of-sample FX interventions show that shocks originating in one currency transmit broadly across FX markets through shared risk exposures.

Keywords: Demand-based asset pricing, Risk factors, Risk exposures, Shock transmission, FX

*We thank Omar Barbiero, Ric Colacito, Hal Cole, Carter Davis, Greg Duffee, Darrell Duffie, Xiang Fang, Thomas Gruenthaler, Sam Hanson, Ben Hébert, Narayana Kocherlakota, Leonid Kogan, Ralph Koijen, Arvind Krishnamurthy, Karen Lewis, Jiacui Li, Wenhao Li, Erik Loualiche, Matteo Maggiori, Stefan Nagel, Daniel Neuhann, Angelo Ranaldo, Dagfinn Rime, Nikolai Roussanov, Davide Tomio, Andrea Vedolin, Adrien Verdelhan, Moto Yogo, and Tony Zhang for helpful discussions. We also thank conference and seminar participants for comments. We gratefully acknowledge funding from the Wharton Dean's Research Fund, the Jacobs Levy Equity Management Center for Quantitative Financial Research at the Wharton School of the University of Pennsylvania, and the General Research Support Fund at Johns Hopkins Carey Business School. This paper was previously circulated under the title "Intermediary Elasticity and Limited Risk-Bearing Capacity." All remaining errors are our own.

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

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

1 Introduction

A defining feature of modern financial markets is their tight interlinkages. Shocks originating in one part of the market can rapidly propagate to many others (e.g., Allen and Gale, 2000; Pavlova and Rigobon, 2008). One important class of shocks is demand shocks. Although demand shocks need not reflect changes in fundamentals, they can move asset prices substantially (e.g., Lee, Shleifer, and Thaler, 1991; Froot and Ramadorai, 2008; Koijen and Yogo, 2019). In an interconnected market, these effects rarely remain confined to the asset initially traded. For example, a currency intervention may move not only the targeted exchange rate but also many other currencies. Which currencies react, and by how much? Answering these questions is essential for effective policy design and response.

We study this question using a novel dataset of daily FX trades between customers and banks from 2012 to 2023 across 17 currencies. The data directly measure the demand imbalances that banks must absorb and allow us to trace how uninformed demand affects exchange rates. Our approach proceeds in three steps. First, we show that cross-currency demand propagation can be decomposed into factor-level repricing and currencies' exposures to common risks. Second, we identify three traded risk factors that account for most of the non-diversifiable risk borne in FX trading and that plausibly exhibit no cross-factor price effects, making the decomposition empirically tractable. Third, we estimate factor-level price sensitivities using sovereign bond auction announcements as instruments for non-informational demand shocks and find that the FX market is highly elastic. Together, these ingredients allow us to recover the cross-currency propagation matrix and validate the resulting estimates using out-of-sample FX interventions.

The empirical challenge is that demand propagation is high-dimensional. Directly estimating all the spillovers across the 17 currencies in our sample is difficult because each exchange rate reflects many demand shocks at once, and clean currency-specific demand shocks are scarce. We reduce this dimensionality using a key asset-pricing insight: banks do not absorb customer demand currency by currency in isolation, but through portfolios of currencies that represent non-diversifiable FX risks, i.e., risk factors. A demand shock to one currency changes banks' exposure to these risk factors; factor prices move; and those price changes transmit to other currencies through their factor exposures.

However, reducing the problem to a small number of factors is not sufficient on its own.

If the price of one factor responds to demand for other factors, or to omitted sources of risk, then estimating factor-level price sensitivities still requires demand shocks that isolate each factor from all other relevant risks. Such factor-specific shocks are empirically difficult to find. We therefore construct factors that plausibly have no cross-factor price spillovers. For these factors, even if there is concurrent customer demand for several factors, the price of each factor responds only to its own demand, allowing us to estimate price sensitivities factor by factor. Our theory shows that these factors can be recovered as currency portfolios with no common variation in both exchange-rate returns and customer-bank flows. We call them "traded risk factors": they are "risk factors" because their returns capture non-diversifiable FX risks, and they are "traded" because their flows capture the customer trading imbalances banks absorb.

The traded risk factors are economically interpretable, stable over time, and quantitatively important. The top three account for about 90% of the non-diversifiable risk banks bear when absorbing customer demand imbalances, confirming that customer demand is absorbed through a small number of non-diversifiable risks. The first two factors resemble the well-known Dollar and Carry factors, but are recovered without using ex ante currency characteristics. The Carry factor, for example, emerges without first sorting currencies by interest rates, as in Lustig, Roussanov, and Verdelhan (2011). The third factor, which we call the Euro-Yen Residual, emerges only when returns and flows are analyzed jointly. It captures the risk banks bear when accommodating active customer trading between the euro area and Japan after Dollar and Carry exposures are hedged. Over the sample, banks sold about $1 trillion of Dollar exposure to customers, accumulated about $0.8 trillion of Carry exposure, and maintained persistently positive exposure to Euro-Yen Residual.

Because the traded risk factors are constructed to have no cross-factor spillovers, we estimate each factor's price sensitivity separately. To do so, we use the week-ahead announcements of sovereign debt auction amounts in eight advanced economies as instruments for customer demand. These announcements are relevant because sovereign auctions attract foreign investors who must acquire local currencies to participate. They are also plausibly exogenous because we use announced offering amounts rather than auction outcomes or realized purchases, thereby exploiting predetermined variation in currency-conversion needs rather than realized demand. Moreover, sovereign issuance is heavily forward guided. For example, in the United States, the Treasury Borrowing Advisory Committee issues recommendations for auction sizes two quarters in advance, and subsequent announcements rarely deviate from those recommendations (Rigon, 2024). As a result, announced offering amounts are unlikely to contain substantial new information about fundamentals or short-term FX market conditions. Consistent with this interpretation, the estimated price effects are temporary and largely reverse within a month.

The estimated factor price sensitivities reveal a highly elastic FX market. A $1 billion demand shock moves factor prices by about 4 basis points for the Dollar factor, 6 basis points for the Carry factor, and 24 basis points for the Euro-Yen Residual. For the Dollar factor, $1 billion corresponds to about 5 basis points of estimated bank dollar positions, implying a demand elasticity of about 1.5. The Dollar exchange rate is therefore substantially more elastic than the aggregate U.S. equity market, where Gabaix and Koijen (2021) estimate an elasticity of about 0.2, but less elastic than the aggregate U.S. Treasury market, where Jansen, Li, and Schmid (2024) estimate an elasticity of about 3.2. One interpretation is that both Dollar FX and U.S. Treasuries are absorbed primarily by large, diversified financial institutions with substantial risk-management capacity, allowing these markets to accommodate demand imbalances more readily than the aggregate equity market.

A key model restriction is that traded risk factors exhibit no cross-factor price spillovers. We test this restriction by examining whether demand shocks for one factor affect the prices of other factors, and find that cross-factor price sensitivities are jointly indistinguishable from zero. This evidence supports treating the traded risk factors as separable and estimating price sensitivities factor by factor.

With no cross-factor spillovers, currency-level propagation can be decomposed into propagation through the traded FX factors. A demand shock to a currency changes the portfolio risks that banks must absorb. Those risk imbalances can be expressed as exposures to the traded FX factors, whose prices change according to the estimated factor price sensitivities. Spillovers to other currencies then depend on their exposures to the same factors, and the effect is obtained by summing across factors. This yields a full cross-currency propagation matrix. Quantitatively, a $1 billion demand shock to one currency moves other exchange rates by up to 9 basis points, indicating substantial propagation across the FX market.

The propagation matrix reveals rich patterns of cross-currency substitution and complementarity. Propagation is strongest when two currencies have similar exposures to the traded risk factors, because demand for one currency raises the prices of risks borne by the other.

Conversely, propagation is weaker when their factor exposures offset, since one currency can help banks hedge the risks generated by demand for the other. For example, the Australian dollar (AUD) and Canadian dollar (CAD) have similar exposures to all three traded risk factors, generating a strong propagation between the two. By contrast, the Japanese yen (JPY) has the opposite Carry exposure relative to AUD and CAD, making it a partial hedge and resulting in much weaker propagation between JPY and those currencies. A more nuanced case is the euro (EUR) and JPY. Both are low-interest-rate currencies and therefore behave as substitutes along the Carry factor, but they lie on opposite sides of the Euro-Yen Residual factor, making them complements along that factor. The two forces offset, leaving their overall cross-currency price sensitivity modest. These findings illustrate that substitution and complementarity are properties of risk exposures rather than currencies themselves: the same pair of currencies can be substitutes for some risks and complements for others.

Finally, we provide an out-of-sample validation of the model-implied cross-currency propagation pattern using spot interventions by the Central Bank of Brazil (BCB). This setting is useful because the interventions in our sample are announced in advance, making the subsequent transactions simply demand shocks to the Brazilian real (BRL). Moreover, the exercise is fully out of sample, as BRL is not used to construct the traded risk factors or to estimate their price sensitivities. For each intervention, we use the estimated factor structure and price sensitivities to predict how a BRL demand shock should affect each sample currency, and compare these predictions with realized exchange-rate responses around the intervention. We find that currencies with larger predicted spillovers exhibit larger realized responses in the predicted direction, providing strong cross-sectional support for the model.

Our paper contributes to the literature on exchange rates by quantifying how demand shocks propagate across currencies through traded risk factors. Beyond conveying information (e.g., Evans and Lyons, 2002; Pasquariello, 2007; Froot and Ramadorai, 2008), customer trading affects exchange rates by changing the non-diversifiable risks that banks must absorb. We recover these risks from joint variation in trading flows and returns. Return comovement across currencies identifies common non-diversifiable risks (Ross, 1976), while flow comovement identifies the portfolios actually traded by customers and absorbed by banks (Hasbrouck and Seppi, 2001). Combining the two lets trading actions reveal which non-diversifiable risks customers trade and banks absorb, providing a trading-based revealed-preference counterpart to the standard approach of constructing factors from conjectured state variables and testing their cross-sectional pricing power.[1] The two most traded risks we recover are the Dollar and Carry factors, the same risks that help price unconditional FX returns in Lustig, Roussanov, and Verdelhan (2011). Our contribution is to show that these are also the dominant risks traded by customers and absorbed by banks. We also uncover a new Euro-Yen Residual factor, highlighting the power of analyzing returns and flows jointly. These findings complement the broader literature on priced currency risks (e.g., Bansal and Dahlquist, 2000; Lustig and Verdelhan, 2007; Hassan and Mano, 2018; Korsaye, Trojani, and Vedolin, 2023) and highlight the role of intermediation frictions in FX pricing (e.g., Gabaix and Maggiori, 2015; Itskhoki and Mukhin, 2021). More broadly, our results shed light on how non-diversifiable risks drive price comovement across currencies, which can also underlie the connection between FX and other asset markets, and in the transmission of monetary policy shocks (e.g., Jiang, Krishnamurthy, and Lustig, 2021; Camanho, Hau, and Rey, 2022; Chernov and Creal, 2023; Gourinchas, Ray, and Vayanos, 2024; Loualiche, Pecora, Somogyi, and Ward, 2024).

Beyond the FX setting, our paper contributes to the broader literature on demand-based asset pricing and inelastic markets (Koijen and Yogo, 2019; Gabaix and Koijen, 2021). A large part of this literature disciplines substitution across assets using investor holdings, asset characteristics, and return covariances (e.g., Kodres and Pritsker, 2002; Pasquariello and Vega, 2015; Koijen and Yogo, 2020; Vayanos and Vila, 2021; Bretscher, Schmid, Sen, and Sharma, 2022; Chaudhary, Fu, and Li, 2023; Davis, Kargar, and Li, 2023; Greenwood, Hanson, and Vayanos, 2023; Jansen, Li, and Schmid, 2024; Jiang, Richmond, and Zhang, 2024; Chaudhry and Davis, 2026). Our contribution is to use the classic factor-versus-idiosyncratic-risk distinction to discipline demand propagation. With N assets, the unrestricted object is an $N \times N$ price-sensitivity matrix, whereby the demand for one asset can move the prices of all its substitutes and complements. Classic asset pricing suggests a risk-based reduction of this object: assets are linked through shared exposures to non-diversifiable risks, while idiosyncratic components primarily generate asset-specific price effects (Markowitz, 1952; Ross, 1976). We apply this logic to demand propagation. Moving from N assets to $K < N$ factors reduces dimensionality, but a generic factor model still allows demand for one factor

[1] For example, Fama and French (1993) identify size and value as key state variables for expected returns, sort stocks by these variables to build the size and value factors, and then show that these factors price the cross-section of expected returns.

to reprice other factors, leaving a $K \times K$ cross-factor price-sensitivity matrix, as emphasized by recent work on cross-asset causal inference and substitution (Haddad, He, Huebner, Kondor, and Loualiche, 2025). We instead construct traded risk factors that are orthogonal in both returns and trading flows and find that cross-factor price spillovers are jointly indistinguishable from zero. The resulting framework reduces the demand-propagation object to K factor-level price sensitivities and currencies' exposures to traded risks.

More generally, our paper contributes to the intermediary asset-pricing and market microstructure literatures, which emphasize limited risk-bearing or balance-sheet capacity as a driver of price responses to demand shocks (e.g., Ho and Stoll, 1981; Grossman and Miller, 1988; Gabaix and Maggiori, 2015; He and Krishnamurthy, 2017; Kondor and Vayanos, 2019; Haddad and Muir, 2021; Du, Hébert, and Huber, 2023; Du, Hébert, and Li, 2023). We share this focus on constrained risk-bearing, but argue that banks' constraints are priced at the level of traded, non-diversifiable risks rather than individual currencies. This perspective is consistent with the foundational insights of Markowitz (1952), Sharpe (1964), and Lintner (1965): what matters for price response is not only the currency being traded, but the risk that banks must absorb.

The next section presents our empirical framework. Section 3 introduces the data. Section 4 constructs the traded risk factors. Section 5 estimates their price sensitivities to demand using sovereign debt auction instruments. Section 6 maps these factor-level estimates into cross-currency demand propagation and provides an out-of-sample validation of the resulting propagation pattern. Section 7 concludes.

2 Empirical Framework

This section provides the framework for quantifying demand propagation in the data. Section 2.1 presents a general bank-customer trading model where equilibrium returns and flows are jointly determined by banks' and customers' willingness to trade. It yields a currency-level supply-demand system in which customer flow into one currency can affect returns in all currencies. The key object is the price sensitivity $\lambda_{n,m}$: the effect of customer flow into currency m on the return of currency n, holding fixed banks' valuation shocks and customer flows into other currencies.

Estimating $\lambda_{n,m}$ directly is difficult because FX demand shocks often move multiple currency flows simultaneously, causing currency-level IV regressions to conflate the direct effect of one currency's flow with the effects of others. The rest of the section addresses this problem by identifying the price sensitivity not of individual currencies but of portfolios of currencies. Section 2.2 first develops a simple benchmark that clarifies why cross impacts disappear for portfolios that experience neither common bank-side nor customer-side shocks. Section 2.3 then constructs such portfolios from return and flow data; we call them traded risk factors. Finally, Section 2.4 shows how factor-level price sensitivities map back to currencies' through no-arbitrage and currency risk exposures. The model therefore delivers a full currency-level demand-propagation matrix even though cross impacts vanish at the traded risk factor level.

2.1 Currency-Level Equilibrium and Price Sensitivities

We begin with an equilibrium model of currency trading between banks and customers. The U.S. dollar, indexed by 0, is the numeraire, and there are N foreign currencies indexed by $n = 1, \ldots, N$. Let $r_{n,t}$ denote the equilibrium excess return of currency n against the U.S. dollar in period t, and let $\Delta Q_{n,t}$ denote the equilibrium customer net flow into currency n. To circumvent the large volume of vehicle-currency trading, we collapse pairwise FX trades into signed flows for non-dollar currencies relative to the U.S. dollar numeraire. For example, if a customer buys euros and sells yen to a bank, the trade is recorded as positive euro flow and negative yen flow. Thus $\Delta Q_{n,t} > 0$ means that customers are net buyers of currency n from banks; banks take the opposite side, so their induced position is $-\Delta Q_{n,t}$.

Appendix A.1 derives the equations below from a general bank-customer equilibrium model. Banks and customers choose currency positions to maximize expected utility; their wealth processes, pre-existing exposures, hedging needs, constraints, and utility functions may differ, and no specific functional form is imposed. The resulting linearized first-order conditions of banks and customers are:

$$r_{n,t} = r^B_{n,t} + \lambda_{n,n}\Delta Q_{n,t} + \sum_{m\neq n} \lambda_{n,m}\Delta Q_{m,t}, \tag{1}$$

$$r_{n,t} = r^C_{n,t} - \theta_{n,n}\Delta Q_{n,t} - \sum_{m\neq n} \theta_{n,m}\Delta Q_{m,t}. \tag{2}$$

These two equations express the same equilibrium return from the two sides of the market.

From the bank side, the return equals $r^B_{n,t}$, banks' valuation of currency n when customer flow is zero, plus the compensation banks require for absorbing customer flows across all currencies. The customer-side equation is analogous, with the opposite sign because flows are measured from the customer side. The valuation shocks $r^B_{n,t}$ and $r^C_{n,t}$ are latent, while $r_{n,t}$ and $\Delta Q_{n,t}$ are observed equilibrium outcomes.

The key structural objects are the slope matrices $\{\lambda_{n,m}\}$ and $\{\theta_{n,m}\}$. The bank-side slopes $\{\lambda_{n,m}\}$ measure how customer flow into currency m changes the return of currency n, holding fixed banks' valuation shock $r^B_{n,t}$. The diagonal terms capture own-currency price sensitivities, while the off-diagonal terms capture demand propagation across currencies. We focus empirically on λ because many important shocks in FX markets shift customer-side FX flows, including bond-issuance-induced flows, benchmark rebalancing, and policy-motivated interventions. The customer-side coefficients $\theta_{n,m}$ are analogous, but identifying them would require bank-side shocks.

Directly estimating $\lambda_{n,m}$ from (1) is difficult even with instruments. In the ideal case, we have, for each currency m, an instrument $z_{m,t}$ that is uncorrelated with bank-side valuation shocks, shifts customer flow into currency m, and does not move customer flows into other currencies:

$$\text{cov}(z_{m,t}, r^B_{n,t}) = 0 \text{ for all } n, \quad \text{cov}(z_{m,t}, \Delta Q_{m,t}) \neq 0, \quad \text{cov}(z_{m,t}, \Delta Q_{j,t}) = 0 \text{ for all } j \neq m. \tag{3}$$

The first two restrictions are the standard IV exclusion and relevance conditions. The last restriction is the additional requirement in a cross-currency setting, whereby the instrument must isolate customer flow into currency m.

This additional requirement is hard to satisfy because FX demand shocks usually affect many currencies at once. For example, a U.K. sovereign bond issuance may raise demand for sterling, but foreign investors can fund the purchase from dollars, euros, yen, or other currencies, generating flows in several currencies at once. Similarly, benchmark rebalancing and policy-motivated interventions change investors' currency portfolios, not just demand for a single currency. Against this reality, instrumenting only $\Delta Q_{m,t}$ would then mix the direct effect of currency m's flow with the effects of flows to other currencies.

This motivates us to directly work with portfolios of currencies. But an arbitrary set of portfolios does not solve the problem, as customer flow into one portfolio can still affect the return of every other portfolio. We need special portfolios for which cross impacts disappear, so that the system can be estimated one portfolio at a time. To build intuition, the next subsection first studies the benchmark in which currencies are counterfactually separable; the following subsection then constructs the relevant separable portfolios from equilibrium return and flow data.

2.2 Benchmark: Separable Currencies

Consider first the benchmark in which each pair of distinct currencies has no common bank-side valuation shock and no common customer-side valuation shock:

$$\operatorname{cov}(r^B_{n,t}, r^B_{m,t}) = 0 \quad \text{and} \quad \operatorname{cov}(r^C_{n,t}, r^C_{m,t}) = 0, \qquad n \neq m. \tag{4}$$

This condition defines the separable-currency benchmark; it is not meant to describe reality.

In this separable benchmark, it is plausible that valuation shocks to one currency do not systematically covary with the equilibrium flow or return of another currency:

$$\operatorname{cov}\left(\begin{pmatrix} r^B_{n,t} \\ r^C_{n,t} \end{pmatrix}, \begin{pmatrix} \Delta Q_{m,t} \\ r_{m,t} \end{pmatrix} \right) = 0_{2\times 2}, \qquad n \neq m. \tag{5}$$

We call the restriction in (5) *Structural Orthogonality.*

The economic rationale is that cross-equilibrium covariance should require a common channel. If currency n's valuation shock systematically covaries with currency m's equilibrium flow or return, then the two currencies must be linked through the bank side, the customer side, or both, such as through a common bank constraint, a common customer trading motive, common information, or shared risk exposure. Once bank-side valuation shocks and customer-side valuation shocks are separately orthogonal, as in (4), it is plausible that no such channel remains.

Structural Orthogonality eliminates the cross terms in the supply-demand system. Appendix A.2 provides a proof.

PROPOSITION 1. *In the separable currency benchmark satisfying Structural Orthogonality, the supply-demand system in* (1) *and* (2) *has no cross impacts. That is, for any two* *distinct currencies $n \neq m$,*

$$\lambda_{n,m} = 0 \qquad \textit{and} \qquad \theta_{n,m} = 0. \tag{6}$$

Thus, in the separable benchmark, the additional currency-level restriction in (3) is no longer needed. Even if an instrument for currency n also moves flows in other currencies, those flows do not affect $r_{n,t}$ once cross impacts vanish. Identification of $\lambda_{n,n}$ therefore only requires customer-flow variation in $\Delta Q_{n,t}$ that is orthogonal to the bank-side valuation shock $r^B_{n,t}$.

The separable-currency benchmark in (4) is stated in terms of latent valuation shocks. The next proposition shows that the same benchmark can be expressed using observed returns and flows. Appendix A.3 provides a proof.

PROPOSITION 2. *Under Structural Orthogonality, for any two distinct currencies $n \neq m$,*

$$\operatorname{cov}(r^B_{n,t}, r^B_{m,t}) = 0 \quad \textit{and} \quad \operatorname{cov}(r^C_{n,t}, r^C_{m,t}) = 0 \tag{7}$$

if and only if

$$\operatorname{cov}(r_{n,t}, r_{m,t}) = 0 \quad \textit{and} \quad \operatorname{cov}(\Delta Q_{n,t}, \Delta Q_{m,t}) = 0. \tag{8}$$

Proposition 2 tells us what to look for in the data. Although bank- and customer-side valuation shocks are latent, returns and customer flows are observed. As raw currencies generally have correlated returns and flows, we next construct the empirical counterpart of the separable benchmark: currency portfolios whose observed returns and flows satisfy (8).

2.3 Traded Risk Factors

We call portfolios whose observed returns and flows satisfy (8) *traded risk factors*. They are "risk factors" because their returns capture non-diversifiable FX risks, and they are "traded" because their flows capture the customer trading imbalances banks absorb.

A factor k is a zero-investment long-short currency portfolio with weights $\{w_{n,k}\}_{n=0}^N$, where the U.S. dollar weight satisfies $w_{0,k} = -\sum_{n=1}^N w_{n,k}$. Since all currency returns are measured relative to the U.S. dollar numeraire, the factor return is

$$r_{k,t}^{\text{factor}} = \sum_{n=1}^{N} w_{n,k} r_{n,t}. \tag{9}$$

The factor flow is obtained from aggregating currency flows by factor exposures. Specifically, let $\beta_{n,k}$ denote currency n's beta with respect to factor k,[2] then customer flow into factor k is

$$\Delta Q_{k,t}^{\text{factor}} = \sum_{n=1}^{N} \beta_{n,k} \Delta Q_{n,t}. \tag{10}$$

As one dollar of customer flow into currency n creates $\beta_{n,k}$ dollars of exposure to factor k, summing these beta-scaled exposures across currencies gives the net customer flow into factor k. Thus $\Delta Q_{k,t}^{\text{factor}}$ measures the factor risk exposure embedded in customer-bank trading (Hull, 2022, Section 3.6).

The following proposition constructs the traded risk factors. Appendix A.4 provides a proof.

PROPOSITION 3. *There exists an ordered set of traded risk factors, defined by currency-portfolio weights $\{w_{n,k}\}$, such that the associated factor returns and factor flows satisfy:*

1. *Uncorrelated factor returns:*

$$\operatorname{cov}\left(r_{k,t}^{\text{factor}}, r_{j,t}^{\text{factor}}\right) = 0, \qquad k \neq j. \tag{11}$$

2. *Uncorrelated factor flows:*

$$\operatorname{cov}\left(\Delta Q_{k,t}^{\text{factor}}, \Delta Q_{j,t}^{\text{factor}}\right) = 0, \qquad k \neq j. \tag{12}$$

The factors are ordered by the amount of trading-induced non-diversifiable risk they capture:

$$\operatorname{var}\left(r_{k,t}^{\text{factor}}\right) \operatorname{var}\left(\Delta Q_{k,t}^{\text{factor}}\right). \tag{13}$$

[2] Formally, let $\mathbf{r}_t^{\text{factor}} := (r_{1,t}^{\text{factor}}, \ldots, r_{K,t}^{\text{factor}})^\top$. For each currency n, define $\boldsymbol{\beta}_n := \operatorname{var}(\mathbf{r}_t^{\text{factor}})^{-1} \operatorname{cov}(\mathbf{r}_t^{\text{factor}}, r_{n,t})$, and let $\beta_{n,k}$ be the k-th element of $\boldsymbol{\beta}_n$.

The factors are unique up to sign and normalization, and up to rotations within eigenspaces associated with repeated eigenvalues.

With the traded risk factors in hand, the separable system can be written factor by factor:

$$r_{k,t}^{\text{factor}} = r_{k,t}^{B,\text{factor}} + \lambda_k^{\text{factor}} \Delta Q_{k,t}^{\text{factor}}, \qquad r_{k,t}^{\text{factor}} = r_{k,t}^{C,\text{factor}} - \theta_k^{\text{factor}} \Delta Q_{k,t}^{\text{factor}}. \tag{14}$$

Thus each traded risk factor has its own price sensitivity to customer demand.

To identify $\lambda_k^{\text{factor}}$, we need an instrument $z_{k,t}^{\text{factor}}$ satisfying the standard IV conditions:

$$\text{cov}\left(z_{k,t}^{\text{factor}}, r_{k,t}^{B,\text{factor}}\right) = 0, \qquad \text{cov}\left(z_{k,t}^{\text{factor}}, \Delta Q_{k,t}^{\text{factor}}\right) \neq 0. \tag{15}$$

Unlike the currency-level condition in (3), the instrument does not need to engender flow in only one single currency or factor because there are no cross-factor impacts, so even if the instrument induces flows across multiple factors, these flows do not enter factor k's return equation.

We use the U.S. dollar as the numeraire only for accounting. Choosing a different numeraire simply re-expresses the same long-short currency portfolios and customer flows in different coordinates. Because the factor construction depends on the joint return-flow covariance structure of these portfolios, the resulting traded risk factors are unchanged. Appendix A.5 provides a proof.

PROPOSITION 4. *The traded risk factors in Proposition 3 are invariant to the choice of numeraire currency.*

2.4 Demand Propagation Through Traded Risk Factors

The coefficient $\lambda_{n,m}$ in (1) measures how customer flow into currency m changes the return of currency n, holding fixed the bank-side valuation shock. The next proposition shows that this currency-level propagation is the sum of repricing effects across traded risk factors. Appendix A.6 provides a proof.

PROPOSITION 5. *The currency-level bank-side price sensitivity is*

$$\lambda_{n,m} = \sum_{k=1}^{K} \frac{\partial \Delta Q_{k,t}^{\text{factor}}}{\partial \Delta Q_{m,t}} \frac{\partial r_{k,t}^{\text{factor}}}{\partial \Delta Q_{k,t}^{\text{factor}}} \frac{\partial r_{n,t}}{\partial r_{k,t}^{\text{factor}}} = \sum_{k=1}^{K} \beta_{m,k} \lambda_k^{\text{factor}} \beta_{n,k}. \tag{16}$$

Equation (16) shows how cross-currency propagation arises even though cross impacts vanish across traded risk factors. A one-dollar customer-flow shock to currency m induces $\beta_{m,k}$ dollars of customer flow into factor k, by the beta aggregation in (10). One dollar of customer flow into factor k moves factor k's return by $\lambda_k^{\text{factor}}$, by the factor-level bank-side equation (14). No-arbitrage then maps this factor repricing back to currencies: a change in factor k's return changes currency n's return by $\beta_{n,k}$.

The formula therefore generates a full and heterogeneous currency-level propagation matrix. Propagation is large when the factor-level price sensitivity $\lambda_k^{\text{factor}}$ is large and when both currencies load strongly on the same factor. Same-sign exposures, $\beta_{m,k}\beta_{n,k} > 0$, amplify propagation because customer flow into the two currencies adds to banks' exposure to the same traded risk. Offsetting exposures, $\beta_{m,k}\beta_{n,k} < 0$, dampen propagation because banks can hedge customer flow in one currency with flow in another. Since the effect sums across factors, two currencies can behave as substitutes along one traded risk and complements along another.

3 Data

This section describes the data used in the empirical analysis. We first introduce the FX trading flow and return data. We then describe the sovereign bond auction data used for identification, and the FX intervention data from the Central Bank of Brazil for validation test.

3.1 Trading Data

Our FX trading data come from the CLS Group (CLS), which provides settlement services for FX trades conducted by its 72 settlement members, primarily large multinational banks.[3]

[3] A list of settlement members can be found at https://www.cls-group.com/communities/settlement-members/.

As the largest single source of FX execution data, CLS covers over 50% of global FX traded volumes.

We use daily aggregate FX order flow data from CLS, which include the total value of buy and sell orders between banks and their customers in 17 currencies from September 2012 to December 2023. The currencies in our sample are: U.S. dollar (USD), Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), Danish kroner (DKK), Euro (EUR), British pound (GBP), Hong Kong dollar (HKD), Israeli shekel (ILS), Japanese yen (JPY), Korean won (KRW), Mexican peso (MXN), Norwegian kroner (NOK), New Zealand dollar (NZD), Swedish kroner (SEK), Singaporean dollar (SGD), and South African rand (ZAR). All trades in our data involve a bank and a customer as counterparties. Banks include bank-affiliated dealers and hedge funds transacting through prime brokers. Customers include asset managers, corporations, and non-bank financials. Our analysis aggregates trading across all banks and all customers, so the empirical object is the market-level net amount bought by customers from banks.

Consistent with the empirical framework in Section 2, we use the customer side of bank-customer trades to construct currency-level flows. To circumvent inflating trading volume from vehicle-currency trading, we aggregate pairwise FX trades into signed flows for each non-dollar currency relative to the U.S. dollar numeraire. Customer purchases are recorded as positive inflows into that currency, while customer sales are recorded as negative inflows; banks take the opposite side of these flows.

To measure the *total* customer FX flow, we are the first to jointly analyze the CLS flows data on FX spot (e.g., Ranaldo and Somogyi, 2021; Roussanov and Wang, 2023) alongside data on FX forwards and swaps. Due to the sometimes negative correlation between flows into spot versus forward and swap, excluding either can mismeasure total customer flow and the corresponding price sensitivity (see Appendix B). The CLS forward and swap data are organized by maturity buckets. We convert these future-settled contracts into comparable spot flow measures by discounting the notional using forward rates.[4,5] Aggregating across

[4]Conceptually, FX swaps should not create net currency positions for banks, as the spot and forward legs offset each other in notional amounts. Empirically, a negligible amount of currency imbalance remains after discounting the forward leg. Our results are effectively unchanged if swaps are excluded.

[5]Specifically, we use the 1-week forward rate for contracts maturing in 1-7 days, the 1-month forward rate for contracts maturing in 8-35 days, the 3-month forward rate for contracts maturing in 36-95 days, and the 1-year forward rate for contracts maturing in more than 96 days. The choice for these rates reflects bucket maturity ranges and forward contract liquidity.

spot, forward, and swap, we construct the USD-valued total daily net customer inflow for each currency.

To align with our instruments, we analyze trading and returns at the weekly frequency. Weekly flows are calculated by summing daily flows from Thursday to the following Wednesday. Our final trading data are a weekly-by-currency panel for 16 non-USD currencies spanning 2012-09-06 to 2023-12-31. Each cell corresponds to $\Delta Q_{n,t}$ in Section 2: the net amount of currency n bought by customers from banks in week t, measured in USD.

3.2 Return Data

We obtain the spot and forward exchange rate data for the 16 non-USD currencies in our sample from Bloomberg. All prices are recorded at the London close. The CLS trading data also follow London business hours.

We define the weekly currency return as the payoff to buying foreign currency forward against the U.S. dollar and converting back at the future spot rate. For currency n from week t to $t+1$, we define $r_{t+1,n} = s_{t+1,n} - f_{t,n}$, where $s_{t,n}$ is the log spot exchange rate and $f_{t,n}$ is the log one-week forward exchange rate (Lustig, Roussanov, and Verdelhan, 2011).[6] Exchange rates are defined as USD per one unit of foreign currency, so a higher s corresponds to USD depreciation.

3.3 Other Data

We collect sovereign bond auction data to construct instruments for FX demand shocks. Specifically, we obtain the auction timing and offering amount of sovereign securities with maturities of at least one year from government websites in the U.S., Australia, France, Germany, Italy, Japan, Mexico, and Sweden. We match the dates of the announcements to the weekly trading and return data described above.

We obtain records of FX interventions by the Central Bank of Brazil (BCB) from its Dados Abertos portal.[7] The BCB conducts several types of interventions, most notably FX

[6]Equivalently, the return can be written as $r_{t+1,n} = s_{t+1,n} - s_{t,n} + i_{t,n} - i_{t,USD} - x_{t,n}$, where $i_{t,n}$ and $i_{t,USD}$ are net risk-free rates and $x_{t,n} = f_{t,n} - s_{t,n} - i_{t,USD} + i_{t,n}$ is the deviation from covered interest-rate parity. Using observed forward rates incorporates these deviations, which helps capture the actual returns faced by banks when absorbing customer flows.

[7]https://www.bcb.gov.br/conteudo/dadosabertos/BCBDepin/historico-atuacoes-mercado-cambio.csv

swaps, which combine a spot intervention with an offsetting forward transaction. Because these operations have limited net effects on spot FX demand, we focus exclusively on spot interventions that are not accompanied by FX swaps.

Finally, we access the Triennial Central Bank Survey of Foreign Exchange and Over-the-Counter Derivatives, conducted by Bank of International Settlement (BIS) from 2002 through 2022 to understand the relevant market size of spot, forward, and swaps.

4 Traded Risk Factors in FX

In this section, we construct traded risk factors in FX using the flow and return data. We show that three factors capture most of the non-diversifiable risk induced by customer FX trading. We then interpret these factors as the Dollar, Carry, and Euro-Yen Residual factors, and describe their cumulative flows and returns over the sample period. Finally, we show that these factors differ from those obtained by standard PCA on returns or flows alone.

4.1 Estimated Traded Risk Factors

Using the procedure detailed in Proposition 3, we construct traded risk factors from weekly net customer flows and returns of 16 non-USD currencies. The three factors that explain the largest share of trading-induced non-diversifiable risk are reported in Table 1. Each column represents a factor, and the entries are the currency weights in this factor. For example, in Factor 1, for every $1 bought, $0.15-worth of CAD and $0.5-worth of EUR are sold.[8] Because the factors are traded portfolios, they place greater weight on widely traded currencies. Notably, seven developed economy currencies, AUD, CAD, CHF, EUR, GBP, JPY, and USD, have consistently high weights across the top three factors; they are highlighted in red.

Of the total trading-induced non-diversifiable risk, the top three traded risk factors individually account for 65%, 16%, and 9%, respectively. Jointly, these three factors explain approximately 90% of the risk banks absorb when accommodating customer trading flows. The traded risk factors are stable over time. Appendix Table A2 shows that the returns

[8] To facilitate comparison, we have scaled such that Factor 1 has a weight of 1 for USD, Factor 2 has all positive weights sum to 1 and all negative weights sum to -1, and Factor 3 has a weight of -1 for JPY. Note that the portfolio weight of USD is the negative sum of the weights of all other currencies.

Table 1: **Top Three Traded Risk Factors**

Currency	Factor 1	Factor 2	Factor 3
AUD	-0.08	0.14	-0.08
CAD	-0.15	0.56	-0.87
CHF	-0.03	-0.07	-0.02
DKK	-0.01	-0.00	0.02
EUR	-0.50	-0.43	1.16
GBP	-0.11	0.18	0.09
HKD	0.00	-0.01	0.02
ILS	0.00	-0.00	0.00
JPY	-0.07	-0.49	-1.00
KRW	-0.01	0.01	-0.01
MXN	-0.01	0.02	-0.03
NOK	-0.01	0.02	-0.01
NZD	-0.01	0.02	-0.01
SEK	-0.01	0.01	-0.01
SGD	-0.01	0.00	0.02
ZAR	-0.01	0.01	-0.01
USD	1.00	0.03	0.74
Var explained	65%	16%	9%

Notes: This table presents the portfolio weights of the top three traded risk factors, constructed following the procedure in Proposition 3. The return and customer flow data for 16 non-USD currencies are weekly from September 2012 to December 2023. The portfolio weight of USD is computed as the negative sum of the weights of all other currencies. The final row reports the share of total trading-induced non-diversifiable risk explained by each factor.

and flows of factors recovered from the full sample are highly correlated with those recovered from the pre-2020 or the post-2020 subsamples: the correlations are near 1 for the first factor and exceed 0.8 for the other two.

4.2 Interpretation of Traded Risk Factors

To interpret the estimated traded risk factors, we compare them with simple portfolios that capture familiar sources of FX risk. Factor 1 in Table 1 assigns negative weights to all non-USD currencies, resembling the Dollar portfolio that shorts non-USD currencies and goes long

USD. As Proposition 4 proves, factors constructed from our procedure are invariant to the choice of numeraire. The Dollar's emergence as the first factor is therefore not mechanical, but reflects its economic role as the world's reserve currency. We construct an interpretive Dollar portfolio that goes long USD and shorts the six most traded currencies, AUD, CAD, CHF, EUR, GBP, and JPY, in equal weights.

Factor 2 has positive weights on high-interest-rate currencies, such as AUD, CAD, and GBP, and negative weights on low-interest-rate currencies, such as JPY, CHF, and EUR. This pattern is consistent with the Carry portfolio that exploits violations of uncovered interest-rate parity. We construct an interpretive Carry portfolio that goes long AUD, CAD, and GBP, and shorts CHF, EUR, and JPY, all in equal weights. Factor 3 features a large positive weight on EUR and a large negative weight on JPY, motivating an interpretive EUR-JPY portfolio that goes long EUR and shorts JPY. This factor may capture large bilateral trading flows between the euro area and Japan. Such flows are not reflected in either the Dollar or the Carry factors, because those factors trade EUR and JPY in the same direction.

These simple interpretive portfolios are economically meaningful but may be correlated. To compare them with the estimated traded risk factors, we apply the same orthogonalization procedure described in Proposition 3. In particular, this procedure transforms the EUR-JPY pair into a Euro-Yen Residual factor that is uncorrelated with the Dollar and Carry factors. Empirically, for every dollar traded in the EUR-JPY pair, 13% of the risk is attributed to the Dollar factor, 25% to the Carry factor, and 62% to the Euro-Yen Residual factor.

The data support this interpretation. Table 2 shows that the returns and flows of the estimated traded risk factors are highly correlated with those of the corresponding interpretive portfolios. The correlations are close to one for both returns and flows across all three factors. Together, the three interpretive portfolios explain about 86% of trading-induced non-diversifiable risk, close to the 90% explained by the estimated traded risk factors. This close alignment supports our interpretation of the estimated factors. In the remainder of the paper, we use the estimated traded risk factors from Table 1 and refer to them as the Dollar, Carry, and Euro-Yen Residual factors.

Table 2: **Estimated Traded Risk Factors and Interpretive Portfolios**

	Factor 1	Factor 2	Factor 3
Return correlation	0.98	0.95	0.92
Flow correlation	1.00	0.99	0.95
Risk explained by interpretive portfolios	63%	15%	8%

Notes: This table compares the estimated traded risk factors in Table 1 with orthogonalized interpretive portfolios constructed from the Dollar, Carry, and EUR-JPY portfolios described in the text. The Return correlation and Flow correlation rows report the correlations between the estimated factors and the corresponding interpretive portfolios. The final row reports the share of total trading-induced non-diversifiable risk explained by each interpretive portfolio.

4.3 Factor Flows and Returns

Having interpreted the three traded risk factors, we next examine their cumulative flows and returns over time. Figure 1 plots cumulative customer flows into the Dollar, Carry, and Euro-Yen Residual factors in Panel (a), and the corresponding cumulative returns in Panel (b).

Panel (a) shows that customers purchased approximately $1 trillion of the Dollar factor from banks during our sample period, primarily after the 2020 COVID crisis. This provision of USD by banks likely reflects USD deposits or wholesale funding made available by dealer-affiliated banks (Du and Huber, 2024). For the Carry factor, customer flows were relatively small in the first part of the sample but turned negative after 2022, indicating customer sales of the Carry factor to banks. As a result, banks accumulated approximately $0.8 trillion in Carry exposure by the end of the sample. Finally, customers sold the Euro-Yen Residual factor before the 2020 COVID crisis and then partially reversed these positions afterward, leaving banks with a net positive position in the Euro-Yen Residual factor over the sample. Because JPY has a negative weight in all three traded risk factors, unwinding JPY positions should not be interpreted mechanically as a change in Carry exposure without analyzing the full factor portfolio. Panel (b) shows that all three traded risk factors have positive cumulative returns over the sample period.

Figure 1: **Cumulative Flows and Returns of Traded Risk Factors**

(a) **Cumulative Flow**

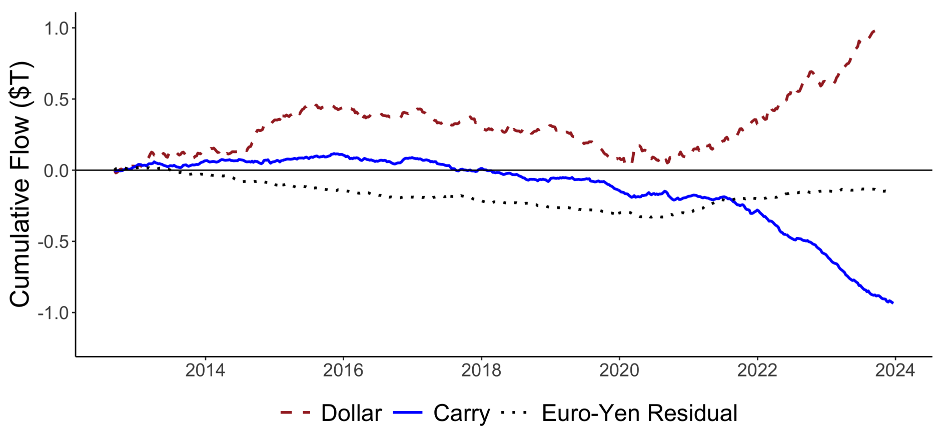

(b) **Cumulative Return**

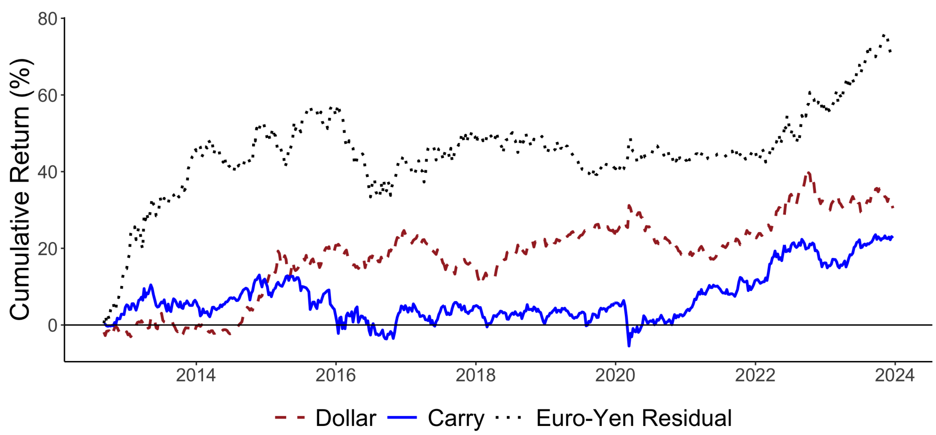

Notes: This figure displays the cumulative flows and returns of the three traded risk factors between September 2012 and December 2023. Flows are measured from the customer perspective: positive flow means that customers buy the factor from banks. For instance, customers bought approximately $1 trillion of the Dollar factor from banks during this period.

4.4 Comparison with PCA on Returns or Flows Alone

The traded risk factors use the joint structure of customer flows and currency returns. To show why both are needed, we compare them with factors obtained from standard PCA applied separately to returns and flows.

Table 3: **Principal Components from FX Returns or Flows Alone**

Currency	Return PCA			Flow PCA		
	PC 1	PC 2	PC 3	PC 1	PC 2	PC 3
AUD	-0.08	0.04	0.27	0.03	0.03	-0.12
CAD	-0.05	0.05	0.32	0.04	1.00	0.06
CHF	-0.05	-0.21	-0.51	0.01	-0.02	0.06
DKK	-0.06	-0.15	-0.12	0.00	-0.00	-0.01
EUR	-0.06	-0.15	-0.13	1.00	-0.03	-0.03
GBP	-0.07	-0.08	0.47	0.02	-0.01	-0.26
HKD	-0.00	0.00	0.00	0.00	-0.02	0.00
ILS	-0.04	-0.03	0.24	0.00	-0.01	0.00
JPY	-0.03	-0.17	-1.00	0.04	-0.06	0.95
KRW	-0.06	0.02	-0.15	0.00	0.01	0.00
MXN	-0.08	0.22	0.71	0.01	0.01	0.00
NOK	-0.10	-0.05	0.72	-0.00	0.01	-0.01
NZD	-0.08	0.01	0.13	0.01	0.01	-0.01
SEK	-0.08	-0.13	0.22	-0.01	0.00	-0.00
SGD	-0.04	-0.03	-0.12	0.01	-0.01	-0.01
ZAR	-0.11	0.29	-1.35	0.01	0.00	-0.01
USD	1.00	0.37	0.29	-1.17	-0.92	-0.62

Notes: The first three columns report portfolio weights for the first three principal components from a return-only PCA, while the next three columns report those from a flow-only PCA. The analysis uses weekly return and customer flow data for 16 non-USD currencies from September 2012 to December 2023. The USD portfolio weight is computed as the negative sum of the weights of all other currencies.

The first three columns of Table 3 show the portfolio weights for the first three principal components of a standard PCA applied to returns.[9] The first factor resembles a Dollar factor, with negative weights on all non-USD currencies. The second factor has some Carry-like features: it assigns large positive weights to high-interest-rate currencies such as ZAR and MXN, and large negative weights to low-interest-rate currencies such as CHF, JPY, and EUR. However, it also assigns very small positive weights to other high-interest-rate

[9]The eigenvectors from a return PCA represent individual currencies' betas to the factors. We convert these betas into portfolio weights using the pseudoinverse of the beta matrix, following the factor-mimicking portfolio approach of Fama and MacBeth (1973).

currencies such as AUD and NZD, and even negative weights to GBP and NOK.[10] The third factor lacks a clear economic interpretation. In contrast, our approach of jointly analyzing flows and returns yields a significant traded risk factor that is unambiguously the Carry and reveals an economically meaningful Euro-Yen Residual factor.

The next three columns of Table 3 report the portfolio weights for the first three principal components of a standard PCA applied to customer flows. The resulting portfolios primarily allocate weight to a single major currency. For instance, the first factor assigns a portfolio weight of -1 to EUR and 0 to all other non-USD currencies, reflecting that EUR/USD is the most actively traded pair. The second and third principal components correspond to the CAD/USD and JPY/USD pairs, respectively. This outcome occurs because the flow PCA identifies portfolios based solely on the largest trading volumes, while overlooking the factor structure in returns.

5 Price Sensitivity of Traded Risk Factors

In this section, we estimate the price sensitivity of traded risk factors to demand shocks. We use sovereign debt auction announcements as instruments for factor flows and estimate how these flows move factor returns. We then test the cross-factor restriction of our empirical framework: a demand shock to one traded risk factor should not move the returns of the other traded risk factors. Finally, we examine the dynamics of the price response. Reversal over time would support a demand-pressure interpretation rather than the alternative that the instruments capture news about fundamentals.

5.1 Empirical Framework and Instruments

Following equation (14) in Section 2, we estimate the price sensitivity of each traded risk factor to its own customer flow. For factor k, we estimate

$$r_{k,t}^{\text{factor}} = \alpha_k + \lambda_k^{\text{factor}} \Delta \widehat{Q}_{k,t}^{\text{factor}} + \epsilon_{k,t}, \tag{17}$$

$$\Delta Q_{k,t}^{\text{factor}} = a_k + \pi_k z_{k,t}^{\text{factor}} + e_{k,t}. \tag{18}$$

[10] Lustig, Roussanov, and Verdelhan (2011) identify the Carry factor from the second principal component after sorting currencies into six portfolios based on interest rate levels.

Here, $r_{k,t}^{\text{factor}}$ is the weekly return of factor k, $\Delta Q_{k,t}^{\text{factor}}$ is the corresponding customer flow, and $z_{k,t}^{\text{factor}}$ is the instrument for this flow. The IV exclusion restriction is

$$\operatorname{cov}\left(z_{k,t}^{\text{factor}}, \epsilon_{k,t}\right) = 0, \tag{19}$$

which requires the instrument to affect factor returns only through the customer flow it induces, and to be otherwise unrelated to the factor return shock. Equation (17) also implies that this price response can be estimated factor by factor: a flow shock to factor k should affect the return of factor k, but not the returns of the other traded risk factors. We test this cross-factor restriction directly in Section 5.3.

We instrument factor flows using sovereign debt auction announcements.[11] Government entities, such as the U.S. Treasury, periodically announce auctions of long-term debt obligations. Foreign investors actively participate in these auctions; for instance, they directly purchased on average 14% of U.S. Treasury notes and bonds sold at auctions between September 2012 and December 2023.[12] When auctions are announced, foreign investors who intend to participate may exchange currencies in advance of purchasing the bonds. Auction announcements therefore generate variation in customer FX flows into the auction currency, making them relevant instruments for factor flows.

The exclusion restriction requires auction announcements to affect factor returns only through the customer flows they induce. Our use of announced offered amounts is central to this restriction. There are two potential sources of information from sovereign auctions that could directly affect exchange rates. The first is information contained in realized investor demand. Although the announcement makes known the amount offered at an auction, it does not reveal how much of the issuance will ultimately be purchased by foreign investors, nor when those investors will convert currencies. Foreign-investor participation is realized later through actual FX trading flows and may itself reflect information, expectations, or other shocks contained in $\epsilon_{k,t}$. For this reason, we use offered amounts rather than auction outcomes, bid-to-cover ratios, or realized foreign purchases. Offered amounts capture predetermined variation in potential currency-conversion needs while avoiding direct use of

[11]We focus on auctions for securities with maturities of longer than one year, as short-term securities are typically bought by domestic investors such as money market funds.

[12]This 14% excludes foreign purchases made indirectly through U.S. investment funds and dealers, so the actual foreign participation may be higher.

realized investor demand.

The second concern is that issuance decisions themselves may convey information. This concern is mitigated by the institutional structure of sovereign debt issuance. Auction schedules and issuance amounts are announced with substantial forward guidance. For example, the U.S. Treasury Borrowing Advisory Committee issues two-quarter-ahead recommendations on debt issuance for upcoming auctions, and subsequent announcements rarely deviate from these recommendations (Rigon, 2024).[13] This institutional structure makes announced offered amounts unlikely to respond to short-term FX market conditions and contain little new information at the time of announcement. Together, these features imply that auction announcements are correlated with factor flows, but should affect factor returns only through the customer flows they induce.

As the traded risk factors place weights on multiple currencies, we construct instruments from a panel of sovereign auction announcements. For each factor, we select sovereign issuers whose currencies are economically representative of the underlying risk factor and whose auction schedules are announced with substantial forward guidance. U.S. Treasury auction announcements instrument flows into the Dollar factor. Australian, Mexican, Japanese, and Swedish sovereign debt auction announcements instrument flows into the Carry factor. Euro-area sovereign debt auction announcements, aggregating German, French, and Italian auctions, instrument flows into the Euro-Yen Residual factor. Within each factor, offered amounts by each sovereign are normalized by their time-series averages so that no single sovereign mechanically dominates the instrument. We then aggregate announcements occurring within the same week of the FX trading data.[14]

5.2 Estimated Price Sensitivity

Table 4 reports the estimated price sensitivity of the Dollar, Carry, and Euro-Yen Residual factors. The dependent variable is the weekly factor return, measured in basis points, and factor flows are measured in billions of U.S. dollars. Therefore, the coefficient on factor flow

[13] As another example, Germany's Finance Agency releases an annual auction calendar each December, specifying target amounts for each auction.

[14] To instrument for factor flows in week t, we use same-week announcements for the Dollar and Carry factors and announcements from weeks $t-1$ and t for the Euro-Yen Residual factor. This longer window accounts for potential delays in auction-induced currency conversion, as Germany, France, and Italy do not allow direct bids from foreign investors.

Table 4: **Estimated Price Sensitivity of Traded Risk Factors**

	OLS			IV		
	Dollar	Carry	Euro-Yen	Dollar	Carry	Euro-Yen
	(1)	(2)	(3)	(4)	(5)	(6)
Factor flow	3.31^{***}	6.23^{***}	13.8^{***}	3.71^{***}	6.20^{**}	23.6^{*}
	(0.301)	(0.713)	(2.19)	(1.43)	(2.54)	(13.8)
1st stage F-stat				32.4	14.9	10.8
Observations	564	564	564	386	188	543

Notes: This table reports the estimation results for $\lambda_k^{\text{factor}}$ in equation (17) for the Dollar, Carry, and Euro-Yen Residual factors. The dependent variable is the weekly factor return, measured in basis points. Factor flows are measured in billions of U.S. dollars, so the coefficient on factor flow is the price response, in basis points, to a $1 billion customer flow into the factor. The IV regressions instrument factor flows using sovereign debt auction announcements. The first-stage F-statistics are heteroscedasticity and autocorrelation consistent F-statistics. The estimation period spans September 2012 to December 2023, excluding the first half of 2020. Newey-West standard errors are reported in parentheses, where the bandwidth is chosen by the Newey and West (1994) selection procedure. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

directly measures the price response, in basis points, to a $1 billion customer flow into the factor.

The estimates are positive for all three traded risk factors. In the IV specifications, a $1 billion customer purchase raises the Dollar factor price by 3.7 basis points, the Carry factor price by 6.2 basis points, and the Euro-Yen Residual factor price by 23.6 basis points. Thus, the Dollar factor has the smallest price sensitivity, the Carry factor is larger, and the Euro-Yen Residual factor is the most sensitive to customer flows. This ordering suggests that banks absorb broad Dollar risk more easily than more specialized risks, such as the Euro-Yen Residual factor, which require larger price adjustments.

The OLS estimates are also positive and statistically significant. For the Dollar and Carry factors, the OLS and IV estimates are similar. For the Euro-Yen Residual factor, the IV estimate is larger but less precisely estimated. One interpretation is that OLS may understate price sensitivity when customer flows are negatively correlated with non-flow return shocks.

For example, if customers tend to buy after adverse valuation shocks or sell after favorable valuation shocks, observed flows partly offset non-flow price movements, attenuating the OLS coefficient. Hence, we use the IV estimates as our main measure of price sensitivity. The first-stage heteroskedasticity and autocorrelation consistent F-statistics are 32.4, 14.9, and 10.8 for the Dollar, Carry, and Euro-Yen Residual factors, respectively, indicating that the auction-announcement instruments are relevant for all three factors.

To compare FX price sensitivities with those in other asset classes, we translate the Dollar-factor estimate into an implied demand elasticity, defined as the demand shock measured as a percentage of estimated bank dollar positions divided by the induced price response. An and Huber (2026) estimate that banks hold approximately \$566 billion of net USD exposure through FX forwards. Using turnover shares to infer the relative size of spot and forward positions implies total bank dollar positions of approximately \$1.85 trillion.[15] A \$1 billion demand shock therefore corresponds to 5.4 basis points of estimated bank dollar positions. Since the estimated Dollar-factor price response is 3.7 basis points, the implied demand elasticity is $5.4/3.7 = 1.46$. This elasticity places Dollar FX between aggregate U.S. equities and U.S. Treasuries: it is substantially larger than the elasticity of about 0.2 implied by the aggregate U.S. equity market estimate of Gabaix and Koijen (2021), but smaller than the aggregate Treasury-market elasticity of about 3.2 implied by Jansen, Li, and Schmid (2024).[16] One interpretation is that demand imbalances in both Dollar FX and U.S. Treasuries are absorbed largely by large, diversified financial institutions with substantial risk-management capacity, allowing these markets to accommodate demand imbalances more readily than the aggregate equity market.

5.3 Cross-Factor Tests

The empirical framework in Section 2 implies that a customer flow shock to one traded risk factor should move that factor's own return, but not the returns of the other traded risk factors. This restriction is empirically testable. We estimate the unrestricted cross-factor IV

[15] According to the BIS Triennial Survey, the average turnover of FX forward to spot is about 0.44 over our sample period; see also Appendix Figure A1.

[16] Gabaix and Koijen (2021) estimate that a 1% demand shock to the aggregate U.S. equity market raises prices by about 5%, implying a demand elasticity of $1/5 = 0.2$. Jansen, Li, and Schmid (2024) estimate an aggregate Treasury-market price multiplier of 0.31. Because the price multiplier is the inverse of market demand, this implies an aggregate Treasury-market elasticity of $1/0.31 = 3.2$.

Table 5: **Tests of Cross-Factor Price Responses**

Tests	F-statistics	p-value
$\lambda_{1,2} = \lambda_{1,3} = 0$	0.65	0.53
$\lambda_{2,1} = \lambda_{2,3} = 0$	0.35	0.70
$\lambda_{3,1} = \lambda_{3,2} = 0$	0.12	0.88
$\lambda_{1,2} = \lambda_{1,3} = \lambda_{2,1} = \lambda_{2,3} = \lambda_{3,1} = \lambda_{3,2} = 0$	0.46	0.83

Notes: This table reports tests of cross-factor coefficients from the unrestricted IV regression in equation (20). The three traded risk factors are ordered as Dollar, Carry, and Euro-Yen Residual. For each return equation k, the first three rows test whether the two coefficients $\lambda_{k,j}^{\text{factor}}$ for $j \neq k$ are jointly zero. The final row tests whether all six cross-factor coefficients are jointly zero. All three factor flows are instrumented jointly using the union of the sovereign auction instruments described in Section 5.1. Test statistics are computed using Newey-West standard errors with the bandwidth selected according to the Newey and West (1994) procedure.

regression

$$r_{k,t}^{\text{factor}} = \alpha_k + \sum_{j=1}^{3} \lambda_{k,j}^{\text{factor}} \Delta \widehat{Q}_{j,t}^{\text{factor}} + \epsilon_{k,t}, \tag{20}$$

where k indexes the factor return and j indexes the factor flow. The own-factor coefficients, $\lambda_{k,k}^{\text{factor}}$, correspond to the price sensitivities estimated in equation (17). The cross-factor coefficients, $\lambda_{k,j}^{\text{factor}}$ for $j \neq k$, capture whether flows into one traded risk factor also move the returns of other factors. The restriction implied by our empirical framework is that these cross-factor coefficients are zero.

Table 5 shows that we do not reject the restriction that cross-factor price responses are zero. In the factor-by-factor tests, the p-values are 0.53, 0.70, and 0.88 for the Dollar, Carry, and Euro-Yen Residual return equations, respectively. The joint test of all six cross-factor restrictions also fails to reject, with a p-value of 0.83. These results support the empirical framework in Section 2 and justify estimating price sensitivities factor by factor in equation (17).

Figure 2: **Reversal of Price Responses**

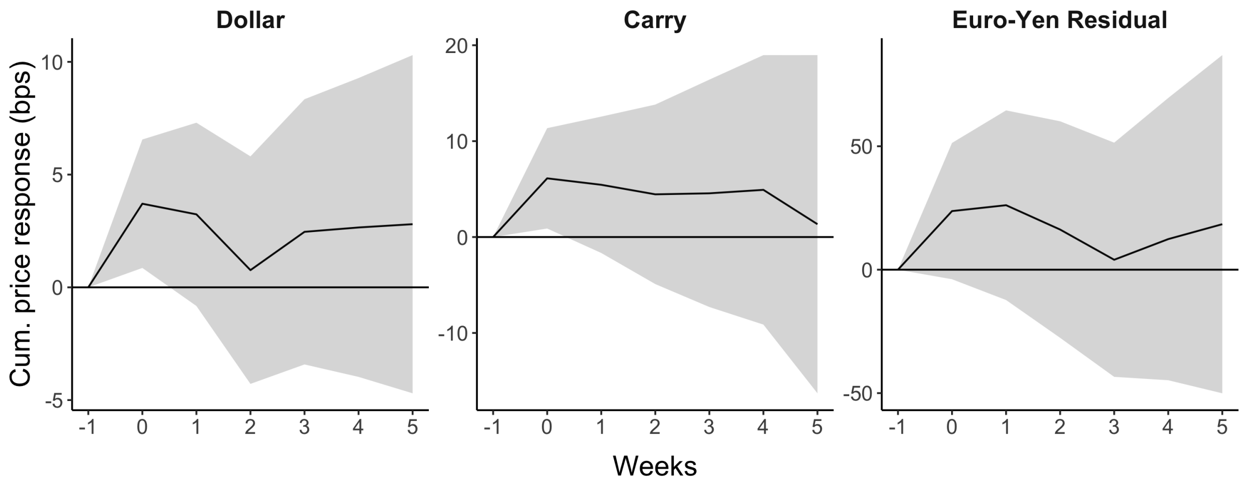

Notes: This figure plots the cumulative price responses $\lambda_{k,h}^{\text{factor}}$ from equation (21) for the Dollar, Carry, and Euro-Yen Residual factors. The point at week -1 is normalized to zero. Week 0 is the contemporaneous response, and weeks 1 through 5 show the cumulative response over subsequent weeks. Responses are measured in basis points per \$1 billion of instrumented customer flow. The shaded area represents the 95% confidence interval based on Newey-West standard errors with the bandwidth selected according to the Newey and West (1994) procedure.

5.4 Reversal of Price Responses

The dynamics of prices provide a further test of the interpretation of the IV estimates. Customer flow shocks can move prices because banks absorb risk, but these price effects should be temporary. By contrast, if the instruments captured news about fundamentals, the induced price response would be persistent.

To examine the dynamics of the price response, for each factor k and horizon h, we estimate

$$\sum_{\ell=0}^{h} r_{k,t+\ell}^{\text{factor}} = \alpha_{k,h} + \lambda_{k,h}^{\text{factor}} \Delta \widehat{Q}_{k,t}^{\text{factor}} + \epsilon_{k,t+h}. \tag{21}$$

The instruments are the same as in the main IV specification in Table 4. The coefficient $\lambda_{k,0}^{\text{factor}}$ captures the contemporaneous price response, while $\lambda_{k,h}^{\text{factor}}$ captures the cumulative price response from the event week through week h.

Figure 2 shows that the price responses are short-lived. All three factors exhibit positive contemporaneous responses, consistent with the estimates in Table 4. Over the following weeks, the cumulative responses move back toward zero and become statistically indistinguishable from zero. This reversal is hard to reconcile with the instruments capturing permanent news about factor valuations. Instead, it supports the interpretation that auction-induced customer flows create temporary price pressure that is absorbed by banks.

Appendix C quantifies the implications of this reversal for factor return variation. Temporary price pressure from customer flows accounts for a sizable share of very short-horizon factor return variation, roughly 10–25% at one week and 5–10% at one month, but fades quickly at longer horizons. Thus, customer-flow price pressure is economically important at short horizons, even though its effect is temporary.

6 Demand Propagation Through Traded Risk Factors

In this section, we map the factor-level price sensitivities estimated above back to individual currencies. This delivers a cross-currency price-sensitivity matrix that shows which currencies move when customers buy a given currency, and by how much. We then use pre-announced FX interventions by the Central Bank of Brazil as an external validation of the model-implied propagation patterns.

6.1 Demand Propagation Across Currencies

For factor-level price sensitivities to generate cross-currency propagation, individual currencies must be exposed to the traded risk factors. Figure 3 shows that they are. We regress currency-level returns on the returns of the Dollar, Carry, and Euro-Yen Residual factors in the time series and plot the marginal R^2 attributed to each factor. These marginal R^2 values are additive because the factor returns are orthogonal by construction. The positive and negative signs in the figure indicate the direction of each currency's beta loading on each factor. Together, the three traded risk factors explain between 53% and 99% of the return variation of the six major non-USD currencies.

The decomposition in Figure 3 provides a bridge from factor-level price sensitivity to currency-level propagation. A customer purchase of one currency changes the factor risks that banks absorb. For example, the three traded risk factors explain 74% of AUD return variation, with the largest component coming from the Dollar factor and smaller components from the Carry and Euro-Yen Residual factors. The signs in the figure show the direction

Figure 3: **Decomposition of Currency Returns Explained by Traded Risk Factors**

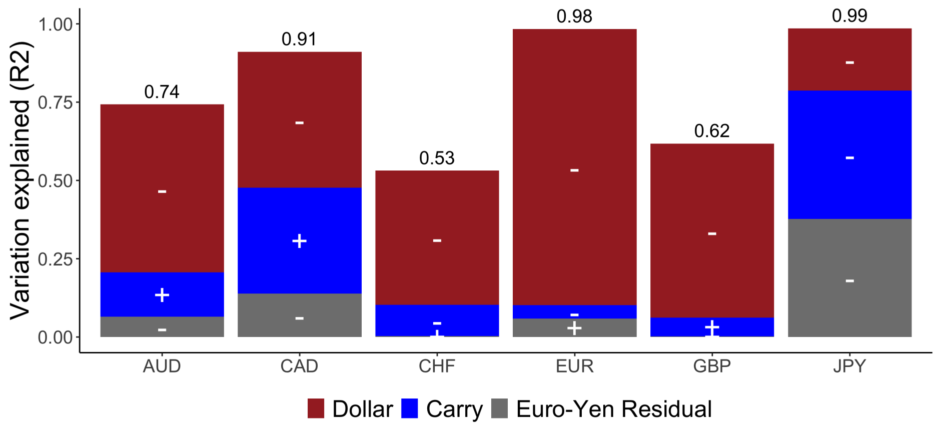

Notes: This figure plots the R^2 of regressing currency-level returns on the returns of the Dollar, Carry, and Euro-Yen Residual factors in the time series. It plots the marginal R^2 values attributed to each factor and labels the total R^2. The positive and negative signs illustrate the direction of the beta loadings.

of these exposures: in factor terms, a customer purchase of AUD corresponds to selling the Dollar and Euro-Yen Residual factors while buying the Carry factor; banks take the opposite side. This same factor repricing also moves other currencies. CAD has the same signs as AUD on all three factors, so CAD rises when the Dollar and Euro-Yen Residual factor prices fall and when the Carry factor price rises.

We quantify this propagation using Proposition 5 and report the cross-currency price sensitivities in Table 6. For clarity, we arrange the six major currencies, AUD, CAD, GBP, CHF, EUR, and JPY, in the upper-left block, followed by the other ten currencies in the sample. Each entry shows the price response in one currency, in basis points, to a $1 billion customer purchase of another currency, holding customer flows in all other currencies fixed. For instance, the AUD-CAD entry of 6.2 means that a $1 billion customer purchase of CAD raises the price of AUD by 6.2 basis points, and symmetrically, a $1 billion customer purchase of AUD raises the price of CAD by 6.2 basis points. The largest estimated cross-currency price sensitivity is 8.7 basis points, between NOK and ZAR.

Table 6 reveals rich patterns of cross-currency propagation. Most cross-currency price sensitivities are positive because all currencies load on the Dollar factor in the same direction, and the Dollar factor is the dominant traded risk factor. However, the magnitude of

Table 6: **Cross-Currency Price Sensitivities Through Traded Risk Factors**

	CAD	GBP	CHF	EUR	JPY	DKK	HKD	ILS	KRW	MXN	NOK	NZD	SEK	SGD	ZAR
AUD	6.2	5.0	2.2	2.8	3.0	2.8	0.1	3.2	4.4	6.0	7.7	6.9	4.7	2.9	8.3
CAD		3.6	0.7	1.0	2.0	1.0	0.1	2.4	3.4	5.1	5.8	5.5	2.9	2.1	6.8
GBP			2.8	4.1	0.2	4.0	0.1	2.6	3.0	4.2	6.3	4.7	5.1	2.2	5.8
CHF				4.6	3.5	4.5	0.0	1.8	1.9	1.0	3.6	2.5	4.3	1.8	2.5
EUR					1.3	5.6	0.1	2.2	2.2	1.9	4.8	3.0	5.5	2.0	3.4
JPY						1.3	0.0	1.5	2.7	0.3	2.2	3.3	1.3	2.2	2.7
DKK							0.1	2.2	2.2	1.9	4.8	3.0	5.5	2.0	3.4
HKD								0.1	0.1	0.1	0.1	0.1	0.1	0.0	0.1
ILS									2.0	2.5	3.7	3.0	2.8	1.5	3.6
KRW										3.3	4.6	4.1	3.1	1.9	4.8
MXN											6.2	5.3	3.7	2.1	6.7
NOK												7.1	6.5	3.3	8.7
NZD													4.6	2.8	7.5
SEK														2.5	5.5
SGD															3.2

Notes: This table uses Proposition 5, the estimated factor-level price sensitivities $\lambda_k^{\text{factor}}$ from Table 4, and the beta loadings of currencies on the traded risk factors to compute currency-level price sensitivities $\lambda_{n,m}$. Each entry reports the price response, in basis points, of the row currency n to a \$1 billion customer purchase of the column currency m, holding customer flows in all other currencies fixed. The model-implied price-sensitivity matrix is symmetric, so we report only the upper half.

propagation varies substantially across currency pairs. Currencies on the long leg of the Carry factor, such as AUD, CAD, and GBP, tend to have larger cross-currency price sensitivities with each other than with currencies on the short leg, such as CHF, EUR, and JPY. For example, AUD and CAD have a cross-currency price sensitivity of 6.2 basis points, while CAD and CHF have a cross-currency price sensitivity of only 0.7 basis points, and GBP and JPY have a cross-currency price sensitivity of only 0.2 basis points. These smaller price sensitivities reflect offsetting exposures to the Carry factor, which makes one currency a partial hedge for the other.

The Euro-Yen Residual factor further shows that the same pair of currencies can be substitutes along one risk and complements along another. EUR and JPY are both low-interest-rate currencies and therefore have similar exposure to the Carry factor. But they lie on opposite sides of the Euro-Yen Residual factor. As a result, their overall cross-currency price sensitivity is only 1.3 basis points, indicating that EUR and JPY are not close substitutes once all traded risk exposures are taken into account.

The table also shows that currencies with small weights in the traded factor portfolios can still have meaningful exposure to the non-diversifiable risks represented by these factors. Take MXN as an example. Its exposure to the traded risk factors implies strong propagation to other high-interest-rate currencies: the cross-currency price sensitivity is 6.0 basis points with AUD, 5.1 basis points with CAD, 6.2 basis points with NOK, 5.3 basis points with NZD, and 6.7 basis points with ZAR. Thus, even currencies that do not receive large portfolio weights in the factor construction can be strongly affected by shocks transmitted through common traded risks.

Finally, the estimated cross-currency price sensitivities provide a useful sanity check through HKD, a currency pegged to USD within a narrow band. We do not use this pegged status in the estimation, but the estimated cross-currency price sensitivities involving HKD are close to zero throughout Table 6. This minimal propagation is exactly what one would expect for a pegged currency: its own customer demand shocks have negligible risk implications for other currencies, and its exchange rate against USD is largely insulated from demand shocks to other currencies.

6.2 Validation Through FX Intervention

The cross-currency price sensitivities in Table 6 are model-implied objects based on the decomposition into traded risk factors. We externally validate the predictions using FX interventions by the Central Bank of Brazil (BCB).

BCB spot interventions provide a useful setting for three reasons. First, we can restrict to interventions that are announced in advance, so information associated with the intervention decision is largely revealed before the transaction takes place. This makes the subsequent intervention closer to a realized demand shock. Second, the Brazilian real (BRL) is not among the currencies used to construct the traded risk factors or estimate their price sensitivities, making the exercise out of sample. Third, the BCB intervention records are publicly available. Over our sample period, there are six pre-announced spot interventions. Consistent with the IV estimation, we exclude the three interventions during the COVID period, leaving us with three test dates.

For each intervention date d, we use the estimated cross-currency price sensitivity $\lambda_{n,BRL}$ to construct the model-predicted propagation from BRL to each sample currency n. Specifically, predicted propagation equals the signed intervention amount multiplied by $\lambda_{n,BRL}$. We then regress realized currency returns around the intervention on this predicted propagation. The unit of observation is a currency-intervention date, and all specifications include intervention-date fixed effects to absorb common market movements and facilitate a cross-sectional test.

Table 7 shows that realized exchange-rate movements line up with model-predicted propagation. The coefficient on predicted propagation is positive and statistically significant for both the one-day and one-week return windows. Thus, currencies with larger predicted exposure to a BRL demand shock move more in the predicted direction. We interpret this exercise as a cross-sectional validation of the propagation pattern rather than a one-for-one calibration test, since intervention amounts may proxy for broader latent BRL demand pressure and event-window returns may contain other currency-specific shocks.

7 Conclusion

This paper quantifies how demand shocks propagate across exchange rates in an interconnected FX market. Using 11 years of customer-bank FX trading flows across 17 currencies,

Table 7: **FX Intervention and Model-Predicted Propagation**

	One-day return (1)	One-week return (2)
Predicted propagation	7.92***	14.90**
	(2.23)	(5.34)
Date FE	Yes	Yes
Observations	48	48

Notes: This table regresses realized exchange-rate returns around non-COVID pre-announced BCB spot interventions on model-predicted propagation. The unit of observation is a currency-intervention date. Predicted propagation is the signed intervention amount, measured in billions of U.S. dollars, multiplied by the estimated cross-currency price sensitivity $\lambda_{n,BRL}$. Column (1) uses the one-day return from the close of the day before the intervention to the close of the intervention day. Column (2) uses the one-week return, defined as the return from two business days before to two business days after the intervention. Both specifications include intervention-date fixed effects. Standard errors are clustered by currency. *$p < .1$; **$p < .05$; ***$p < .01$.

we show that cross-currency propagation can be understood through a small number of traded risk factors. Rather than estimating hundreds of bilateral spillovers directly, we decompose propagation into factor-level repricing and currency-specific risk exposures. The key innovation is to identify traded risk factors from the joint dynamics of prices and quantities and to construct them so that they plausibly exhibit no cross-factor price spillovers.

To estimate factor-level price sensitivities, we introduce sovereign debt auction announcements as instruments for demand shocks. We find that FX markets are highly elastic and that different traded risk factors exhibit markedly different price sensitivities. Combining these estimates with currencies' factor exposures yields a complete cross-currency propagation matrix. The resulting spillovers are economically large: a $1 billion demand shock to one currency can move other exchange rates by up to 9 basis points. We further show that the model's predicted spillovers align closely with realized responses to out-of-sample interventions by the Central Bank of Brazil.

More broadly, our framework links trading quantities and asset prices through common risks (Froot and Ramadorai, 2008; Koijen and Yogo, 2019). As financial markets become increasingly integrated, understanding how common risks transmit demand shocks across assets is essential for measuring market resilience, anticipating spillovers, and evaluating policy interventions.

Bibliography

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

An, Y., and A. W. Huber. 2026. Geoeconomic competition and capital reallocation in global fx funding. Working Paper.

Bansal, R., and M. Dahlquist. 2000. The forward premium puzzle: different tales from developed and emerging economies. *Journal of International Economics* 51:115–44.

Bretscher, L., L. Schmid, I. Sen, and V. Sharma. 2022. Institutional corporate bond pricing. Working Paper.

Camanho, N., H. Hau, and H. Rey. 2022. Global portfolio rebalancing and exchange rates. *The Review of Financial Studies* 35:5228–74.

Chaudhary, M., Z. Fu, and J. Li. 2023. Corporate bond multipliers: Substitutes matter. Working Paper.

Chaudhry, A., and C. Davis. 2026. The origins of the factor zoo: Investors weakly substitute across stocks. Working paper.

Chernov, M., and D. Creal. 2023. International yield curves and currency puzzles. *The Journal of Finance* 78:209–45.

Davis, C., M. Kargar, and J. Li. 2023. Why is asset demand inelastic? Working Paper.

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

Du, W., B. Hébert, and W. Li. 2023. Intermediary balance sheets and the treasury yield curve. *Journal of Financial Economics* 150:103722–.

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

Evans, M., and R. Lyons. 2002. Order flow and exchange rate dynamics. *Journal of Political Economy* 110:170–80.

Fama, E. F., and K. R. French. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33:3–56.

Fama, E. F., and J. D. MacBeth. 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy* 81:607–36.

Froot, K. A., and T. Ramadorai. 2008. Institutional portfolio flows and international investments. *Review of Financial Studies* 21:937–71.

Gabaix, X., and R. Koijen. 2021. In search of the origins of financial fluctuations: the inelastic market hypothesis. Working Paper.

Gabaix, X., and M. Maggiori. 2015. International liquidity and exchange rate dynamics. *Quarterly Journal of Economics* 130:1369–420.

Gourinchas, P.-O., W. Ray, and D. Vayanos. 2024. A preferred-habitat model of term premia, exchange rates, and monetary policy spillovers. Working Paper.

Greenwood, R., S. Hanson, and D. Vayanos. 2023. Supply and demand and the term structure of interest rates. *Annual Review of Financial Economics* 16.

Grossman, S. J., and M. H. Miller. 1988. Liquidity and market structure. *The Journal of Finance* 43:617–33.

Haddad, V., Z. He, P. Huebner, P. Kondor, and E. Loualiche. 2025. Causal inference for asset pricing. Working Paper.

Haddad, V., and T. Muir. 2021. Do intermediaries matter for aggregate asset prices? *The Journal of Finance* 76:2719–61.

Hasbrouck, J., and D. J. Seppi. 2001. Common factors in prices, order flows, and liquidity. *Journal of Financial Economics* 59:383–411.

Hassan, T. A., and R. C. Mano. 2018. Forward and spot exchange rates in a multi-currency world. *The Quarterly Journal of Economics* 134:397–450.

———. 2017. Intermediary asset pricing and the financial crisis. *Annual Review of Financial Economics* 173–97.

Ho, T., and H. R. Stoll. 1981. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics* 9:47–73.

Hull, J. C. 2022. *Options, futures, and other derivatives.* 11 ed. Pearson.

Itskhoki, O., and D. Mukhin. 2021. Exchange rate disconnect in general equilibrium. *Journal of Political Economy* 129:2183–232.

Jansen, K. A., W. Li, and L. Schmid. 2024. Granular treasury demand with arbitrageurs Working Paper.

Jiang, Z., A. Krishnamurthy, and H. Lustig. 2021. Foreign safe asset demand and the dollar exchange rate. *The Journal of Finance* 76:1049–89.

Jiang, Z., R. J. Richmond, and T. Zhang. 2024. A portfolio approach to global imbalances. *The Journal of Finance* 79:2025–76.

Kodres, L. E., and M. Pritsker. 2002. A rational expectations model of financial contagion. *Journal of Finance* 57:769–99.

Koijen, R. S. J., and M. Yogo. 2019. A demand system approach to asset pricing. *Journal of Political Economy* 127:1475–515.

———. 2020. Exchange rates and asset prices in a global demand system. Working Paper.

Kondor, P., and D. Vayanos. 2019. Liquidity risk and the dynamics of arbitrage capital. *Journal of Finance* 74:1139–73.

Korsaye, S. A., F. Trojani, and A. Vedolin. 2023. The global factor structure of exchange rates. *Journal of Financial Economics* 148:21–46.

Lee, C. M., A. Shleifer, and R. H. Thaler. 1991. Investor sentiment and the closed-end fund puzzle. *The Journal of Finance* 46:75–109.

Lintner, J. 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics* 47:13–37.

Loualiche, E., A. R. Pecora, F. Somogyi, and C. Ward. 2024. Monetary policy transmission through the exchange rate factor structure. Working Paper.

Lustig, H., N. Roussanov, and A. Verdelhan. 2011. Common Risk Factors in Currency Markets. *The Review of Financial Studies* 24:3731–77.

Lustig, H., and A. Verdelhan. 2007. The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review* 97:89–117.

Markowitz, H. 1952. Portfolio selection. *Journal of Finance* 7:77–91.

Newey, W. K., and K. D. West. 1994. Automatic lag selection in covariance matrix estimation. *Review of Economic Studies* 61:631–53.

Pasquariello, P. 2007. Informative trading or just costly noise? An analysis of central bank interventions. *Journal of Financial Markets* 10:107–43.

Pasquariello, P., and C. Vega. 2015. Strategic cross-trading in the US stock market. *Review of Finance* 19:229–82.

Pavlova, A., and R. Rigobon. 2008. The role of portfolio constraints in the international propagation of shocks. *The Review of Economic Studies* 75:1215–56.

Ranaldo, A., and F. Somogyi. 2021. Asymmetric information risk in FX markets. *Journal of Financial Economics* 140:391–411.

Rigon, L. 2024. Three essays on US central banking and debt management. Stanford University Ph.D. thesis.

Ross, S. A. 1976. The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13:341–60.

Roussanov, N., and X. Wang. 2023. Following the Fed: Limits of arbitrage and the dollar. Working Paper.

Sharpe, W. F. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance* 19:425–42.

Vayanos, D., and J.-L. Vila. 2021. A preferred-habitat model of the term structure of interest rates. *Econometrica* 89:77–112.

A Derivations and Proofs

This appendix collects the derivations and proofs omitted from the main text.

A.1 Derivation of the Currency-Level Equilibrium System

We derive the currency-level system in (1) and (2). The U.S. dollar is the numeraire. Let $P_{n,t}$ denote the time-t spot price of currency n, measured in U.S. dollars per unit of foreign currency. Let R_F denote the gross one-period U.S. dollar risk-free rate, and let $R_{F,n}$ denote the gross one-period risk-free rate in currency n. The excess return of currency n is

$$r_{n,t} := R_{F,n} \frac{P_{n,t}}{P_{n,t-1}} - R_F. \tag{A1}$$

A signed position q_n in currency n, measured in U.S. dollars at the pre-trade price $P_{n,t-1}$, corresponds to $q_n/P_{n,t-1}$ units of foreign currency. If this position is acquired at the equilibrium spot price $P_{n,t}$, invested at the foreign risk-free rate $R_{F,n}$, and financed at the U.S. dollar risk-free rate R_F, its net payoff at $t+1$ is

$$q_n \left(R_{F,n} \frac{P_{n,t+1}}{P_{n,t-1}} - R_F \frac{P_{n,t}}{P_{n,t-1}} \right). \tag{A2}$$

Using (A1), define the payoff to one dollar of signed position in currency n, expressed as a function of the contemporaneous excess return $r_{n,t}$, by

$$\Pi_{n,t+1}(r_{n,t}) := R_{F,n} \frac{P_{n,t+1}}{P_{n,t-1}} - \frac{R_F}{R_{F,n}} (r_{n,t} + R_F). \tag{A3}$$

This is equivalent to the payoff expression in (A2).

Banks and customers choose currency positions by solving

$$\{-\Delta Q_{n,t}\}_{n=1}^N = \arg \max_{\{q_n\}_{n=1}^N} \mathbb{E}_t \left[u_B \left(W_{B,t+1} + \sum_{n=1}^N q_n \Pi_{n,t+1}(r_{n,t}) \right) \right], \tag{A4}$$

and

$$\{\Delta Q_{n,t}\}_{n=1}^N = \arg \max_{\{q_n\}_{n=1}^N} \mathbb{E}_t \left[u_C \left(W_{C,t+1} + \sum_{n=1}^N q_n \Pi_{n,t+1}(r_{n,t}) \right) \right]. \tag{A5}$$

The wealth processes $W_{B,t+1}$ and $W_{C,t+1}$ may be risky and may differ between banks and customers, capturing pre-existing currency positions, hedging needs, and other exposures. The utility functions u_B and u_C may also differ. We assume only that they are twice continuously differentiable, strictly increasing, and strictly concave.

We prove the bank-side representation (1). The customer-side representation (1) is analogous. For a bank position $\mathbf{q} = (q_1, \ldots, q_N)$, define bank terminal wealth by

$$Y^B_{t+1}(\mathbf{q}, \mathbf{r}_t) := W_{B,t+1} + \sum_{m=1}^{N} q_m \Pi_{m,t+1}(r_{m,t}). \tag{A6}$$

The bank first-order condition for currency n is

$$F_{n,t}(\mathbf{q}, \mathbf{r}_t) := \mathbb{E}_t \left[u'_B\big(Y^B_{t+1}(\mathbf{q}, \mathbf{r}_t)\big)\, \Pi_{n,t+1}(r_{n,t}) \right] = 0. \tag{A7}$$

The no-flow valuation $\mathbf{r}^B_t = (r^B_{1,t}, \ldots, r^B_{N,t})$ is defined by imposing this first-order condition at $\mathbf{q} = \mathbf{0}$. At this point,

$$Y^B_{t+1}(\mathbf{0}, \mathbf{r}_t) = W_{B,t+1}. \tag{A8}$$

Therefore,

$$F_{n,t}(\mathbf{0}, \mathbf{r}^B_t) = \mathbb{E}_t \left[u'_B(W_{B,t+1}) \Pi_{n,t+1}(r^B_{n,t}) \right] = 0. \tag{A9}$$

Substituting (A3) into (A9) gives the bank-side no-flow valuation equation:

$$\mathbb{E}_t \left[u'_B(W_{B,t+1}) \left(R_{F,n} \frac{P_{n,t+1}}{P_{n,t-1}} - \frac{R_F}{R_{F,n}} (r^B_{n,t} + R_F) \right) \right] = 0. \tag{A10}$$

Moreover, define the bank payoff residual at the no-flow valuation by

$$X^B_{n,t+1} := \Pi_{n,t+1}(r^B_{n,t}). \tag{A11}$$

We now apply the implicit function theorem to the system

$$\mathbf{F}_t(\mathbf{q}, \mathbf{r}_t) = \mathbf{0}, \tag{A12}$$

where $\mathbf{F}_t = (F_{1,t}, \ldots, F_{N,t})^\top$. For $n, \ell \in \{1, \ldots, N\}$, the derivative of $F_{n,t}$ with respect to $r_{\ell,t}$

is

$$\frac{\partial F_{n,t}}{\partial r_{\ell,t}} = \mathbb{E}_t \left[u_B''(Y_{t+1}^B) \frac{\partial Y_{t+1}^B}{\partial r_{\ell,t}} \Pi_{n,t+1}(r_{n,t}) + u_B'(Y_{t+1}^B) \frac{\partial \Pi_{n,t+1}(r_{n,t})}{\partial r_{\ell,t}} \right]. \tag{A13}$$

Since

$$\frac{\partial Y_{t+1}^B}{\partial r_{\ell,t}} = q_\ell \left(-\frac{R_F}{R_{F,\ell}} \right), \tag{A14}$$

this term vanishes at the no-flow point $\mathbf{q} = \mathbf{0}$. Also,

$$\frac{\partial \Pi_{n,t+1}(r_{n,t})}{\partial r_{\ell,t}} = -\frac{R_F}{R_{F,n}} \mathbf{1}\{n = \ell\}. \tag{A15}$$

Evaluating (A13) at $(\mathbf{q}, \mathbf{r}_t) = (\mathbf{0}, \mathbf{r}_t^B)$ gives

$$\left. \frac{\partial F_{n,t}}{\partial r_{\ell,t}} \right|_{(\mathbf{0}, \mathbf{r}_t^B)} = -\frac{R_F}{R_{F,n}} \mathbb{E}_t[u_B'(W_{B,t+1})] \mathbf{1}\{n = \ell\}. \tag{A16}$$

Because $u_B' > 0$, the Jacobian of $\mathbf{F}_t$ with respect to $\mathbf{r}_t$ is diagonal with nonzero diagonal entries, hence invertible.

Next, for $n, m \in \{1, \ldots, N\}$,

$$\frac{\partial F_{n,t}}{\partial q_m} = \mathbb{E}_t \left[u_B''(Y_{t+1}^B) \frac{\partial Y_{t+1}^B}{\partial q_m} \Pi_{n,t+1}(r_{n,t}) \right], \tag{A17}$$

and

$$\frac{\partial Y_{t+1}^B}{\partial q_m} = \Pi_{m,t+1}(r_{m,t}). \tag{A18}$$

Evaluating at $(\mathbf{q}, \mathbf{r}_t) = (\mathbf{0}, \mathbf{r}_t^B)$ gives

$$\left. \frac{\partial F_{n,t}}{\partial q_m} \right|_{(\mathbf{0}, \mathbf{r}_t^B)} = \mathbb{E}_t \left[u_B''(W_{B,t+1}) X_{n,t+1}^B X_{m,t+1}^B \right]. \tag{A19}$$

By the implicit function theorem, there exists a differentiable local function $\mathbf{r}_t(\mathbf{q})$ satisfying $\mathbf{F}_t(\mathbf{q}, \mathbf{r}_t(\mathbf{q})) = \mathbf{0}$, with derivative

$$\left. \frac{\partial r_{n,t}}{\partial q_m} \right|_{\mathbf{q}=\mathbf{0}} = \frac{R_{F,n} \mathbb{E}_t \left[u_B''(W_{B,t+1}) X_{n,t+1}^B X_{m,t+1}^B \right]}{R_F \mathbb{E}_t [u_B'(W_{B,t+1})]}. \tag{A20}$$

The bank position is the negative of customer flow, so $q_m = -\Delta Q_{m,t}$. Hence

$$\left.\frac{\partial r_{n,t}}{\partial \Delta Q_{m,t}}\right|_{\Delta \mathbf{Q}_t = \mathbf{0}} = \frac{R_{F,n}\mathbb{E}_t\left[-u_B''(W_{B,t+1})X_{n,t+1}^B X_{m,t+1}^B\right]}{R_F\mathbb{E}_t[u_B'(W_{B,t+1})]} := \lambda_{n,m,t}. \tag{A21}$$

Therefore, the first-order Taylor expansion of $r_{n,t}$ around $\Delta \mathbf{Q}_t = \mathbf{0}$ is

$$r_{n,t} = r_{n,t}^B + \sum_{m=1}^{N} \lambda_{n,m,t}\Delta Q_{m,t} + o(\|\Delta \mathbf{Q}_t\|). \tag{A22}$$

The derivation delivers a state-dependent local sensitivity $\lambda_{n,m,t}$. In the main text and empirical analysis, we impose the time-invariant approximation $\lambda_{n,m,t} \equiv \lambda_{n,m}$, interpreted as the average local price sensitivity over the sample. Dropping the higher-order term then gives

$$r_{n,t} = r_{n,t}^B + \lambda_{n,n}\Delta Q_{n,t} + \sum_{m \neq n} \lambda_{n,m}\Delta Q_{m,t}, \tag{A23}$$

which proves the bank-side equation used in the main text.

Repeating the same argument for the customer problem and imposing $\theta_{n,m,t} \equiv \theta_{n,m}$ gives

$$r_{n,t} = r_{n,t}^C - \theta_{n,n}\Delta Q_{n,t} - \sum_{m \neq n} \theta_{n,m}\Delta Q_{m,t}. \tag{A24}$$

The sign differs because $\Delta Q_{n,t}$ is measured from the customer side: customers' position is $+\Delta Q_{n,t}$, while banks' induced position is $-\Delta Q_{n,t}$.

A.2 Proof of Proposition 1

Suppress the time subscript. Let $\mathbf{r}, \mathbf{r}^B, \mathbf{r}^C, \Delta\mathbf{Q}$ denote the vectors of equilibrium returns, bank-side valuation shocks, customer-side valuation shocks, and customer flows. Let $\mathbf{\Lambda}$ and $\mathbf{\Theta}$ denote the slope matrices in (1) and (2).

Define the normalized variables and coefficients

$$\widetilde{r}_n := \frac{r_n}{R_{F,n}}, \qquad \widetilde{r}_n^B := \frac{r_n^B}{R_{F,n}}, \qquad \widetilde{r}_n^C := \frac{r_n^C}{R_{F,n}}, \tag{A25}$$

and

$$\widetilde{\lambda}_{n,m} := \frac{\lambda_{n,m}}{R_{F,n}}, \qquad \widetilde{\theta}_{n,m} := \frac{\theta_{n,m}}{R_{F,n}}. \tag{A26}$$

The definition of $\lambda_{n,m,t}$ in (A21), together with the time-invariant approximation $\lambda_{n,m,t} \equiv \lambda_{n,m}$, implies that the matrix with entries $\widetilde{\lambda}_{n,m}$ is symmetric; the analogous customer-side expression implies that the matrix with entries $\widetilde{\theta}_{n,m}$ is symmetric. Since $R_{F,n} > 0$, this row normalization does not affect zero restrictions: $\lambda_{n,m} = 0$ if and only if $\widetilde{\lambda}_{n,m} = 0$, and similarly for $\theta_{n,m}$.

In normalized units, the local system is

$$\widetilde{\mathbf{r}} = \widetilde{\mathbf{r}}^B + \widetilde{\mathbf{\Lambda}}\Delta\mathbf{Q}, \qquad \widetilde{\mathbf{r}} = \widetilde{\mathbf{r}}^C - \widetilde{\mathbf{\Theta}}\Delta\mathbf{Q}. \tag{A27}$$

The covariance restrictions in (4) and (5) are unchanged by the positive row normalization in (A25). Hence, in normalized units, $\text{var}(\widetilde{\mathbf{r}}^B)$, $\text{var}(\widetilde{\mathbf{r}}^C)$, $\text{cov}(\widetilde{\mathbf{r}}^B, \Delta\mathbf{Q})$, $\text{cov}(\widetilde{\mathbf{r}}^C, \Delta\mathbf{Q})$, $\text{cov}(\widetilde{\mathbf{r}}^B, \widetilde{\mathbf{r}})$, and $\text{cov}(\widetilde{\mathbf{r}}^C, \widetilde{\mathbf{r}})$ are diagonal.

Define

$$\mathbf{D} := \text{cov}(\widetilde{\mathbf{r}}^B, \Delta\mathbf{Q}), \qquad \mathbf{A} := \text{cov}(\widetilde{\mathbf{r}}^C, \Delta\mathbf{Q}), \qquad \mathbf{\Sigma}_Q := \text{var}(\Delta\mathbf{Q}), \tag{A28}$$

and write the diagonal entries of $\mathbf{D}$ and $\mathbf{A}$ as d_n and a_n. Let

$$\mathbf{G} := \widetilde{\mathbf{\Lambda}} + \widetilde{\mathbf{\Theta}}. \tag{A29}$$

For the nondegenerate directions we keep in the model, $\mathbf{G}$ and $\mathbf{\Sigma}_Q$ are nonsingular.

From the bank-side equation in (A27),

$$\text{cov}(\widetilde{\mathbf{r}}^B, \widetilde{\mathbf{r}}) = \text{var}(\widetilde{\mathbf{r}}^B) + \mathbf{D}\widetilde{\mathbf{\Lambda}}. \tag{A30}$$

The left-hand side and $\text{var}(\widetilde{\mathbf{r}}^B)$ are diagonal, so $\mathbf{D}\widetilde{\mathbf{\Lambda}}$ is diagonal. Therefore, for $n \neq m$,

$$d_n \widetilde{\lambda}_{n,m} = 0. \tag{A31}$$

Similarly, from the customer-side equation in (A27),

$$\operatorname{cov}(\widetilde{\mathbf{r}}^C, \widetilde{\mathbf{r}}) = \operatorname{var}(\widetilde{\mathbf{r}}^C) - \mathbf{A}\widetilde{\boldsymbol{\Theta}}. \tag{A32}$$

The left-hand side and $\operatorname{var}(\widetilde{\mathbf{r}}^C)$ are diagonal, so $\mathbf{A}\widetilde{\boldsymbol{\Theta}}$ is diagonal. Therefore, for $n \neq m$,

$$a_n \widetilde{\theta}_{n,m} = 0. \tag{A33}$$

Combining the two equations in (A27) gives

$$\widetilde{\mathbf{r}}^C - \widetilde{\mathbf{r}}^B = \mathbf{G}\Delta\mathbf{Q}. \tag{A34}$$

Taking covariance with $\Delta\mathbf{Q}$ yields

$$\mathbf{A} - \mathbf{D} = \mathbf{G}\boldsymbol{\Sigma}_Q. \tag{A35}$$

Since $\mathbf{G}$ and $\boldsymbol{\Sigma}_Q$ are nonsingular, $\mathbf{A}-\mathbf{D}$ is nonsingular. Since $\mathbf{A}-\mathbf{D}$ is diagonal, this implies

$$a_n - d_n \neq 0 \qquad \text{for every } n. \tag{A36}$$

Now use $\widetilde{\mathbf{r}}^C = \widetilde{\mathbf{r}}^B + \mathbf{G}\Delta\mathbf{Q}$. Since $\mathbf{G}$ is symmetric,

$$\operatorname{var}(\widetilde{\mathbf{r}}^C) = \operatorname{var}(\widetilde{\mathbf{r}}^B) + \mathbf{DG} + \mathbf{GD} + \mathbf{G}\boldsymbol{\Sigma}_Q\mathbf{G}. \tag{A37}$$

Using (A35), the last term is $\mathbf{G}\boldsymbol{\Sigma}_Q\mathbf{G} = (\mathbf{A}-\mathbf{D})\mathbf{G}$. Since $\operatorname{var}(\widetilde{\mathbf{r}}^B)$ and $\operatorname{var}(\widetilde{\mathbf{r}}^C)$ are diagonal, the (n, m) off-diagonal entry of (A37) gives, for $n \neq m$,

$$(a_n + d_m)g_{n,m} = 0, \qquad g_{n,m} := \widetilde{\lambda}_{n,m} + \widetilde{\theta}_{n,m}. \tag{A38}$$

Fix $n \neq m$. Suppose, toward contradiction, that $\widetilde{\lambda}_{n,m} \neq 0$. Since $\widetilde{\boldsymbol{\Lambda}}$ is symmetric, $\widetilde{\lambda}_{m,n} \neq 0$. By (A31), $d_n = 0$ and $d_m = 0$. If $\widetilde{\theta}_{n,m} \neq 0$, symmetry of $\widetilde{\boldsymbol{\Theta}}$ and (A33) imply $a_n = 0$ and $a_m = 0$, contradicting (A36). Hence $\widetilde{\theta}_{n,m} = 0$. Then $g_{n,m} = \widetilde{\lambda}_{n,m} \neq 0$, so (A38) implies $a_n + d_m = 0$. Since $d_m = 0$, this gives $a_n = 0$, contradicting $a_n - d_n \neq 0$ because $d_n = 0$. Therefore $\widetilde{\lambda}_{n,m} = 0$.

Now suppose, toward contradiction, that $\widetilde{\theta}_{n,m} \neq 0$. Since $\widetilde{\boldsymbol{\Theta}}$ is symmetric, $\widetilde{\theta}_{m,n} \neq 0$. By (A33), $a_n = 0$ and $a_m = 0$. We already proved $\widetilde{\lambda}_{n,m} = 0$, so $g_{n,m} = \widetilde{\theta}_{n,m} \neq 0$. Equation (A38) implies $a_n + d_m = 0$. Since $a_n = 0$, this gives $d_m = 0$, contradicting $a_m - d_m \neq 0$. Hence $\widetilde{\theta}_{n,m} = 0$.

Thus $\widetilde{\lambda}_{n,m} = 0$ and $\widetilde{\theta}_{n,m} = 0$ for all $n \neq m$. Since $R_{F,n} > 0$, the original coefficients also satisfy $\lambda_{n,m} = 0$ and $\theta_{n,m} = 0$ for all $n \neq m$.

A.3 Proof of Proposition 2

Suppress the time subscript. By Proposition 1, the system is diagonal:

$$r_n = r_n^B + \lambda_n \Delta Q_n, \qquad r_n = r_n^C - \theta_n \Delta Q_n. \tag{A39}$$

Let $g_n := \lambda_n + \theta_n$, and assume $g_n \neq 0$. Combining the two equations gives

$$\Delta Q_n = \frac{r_n^C - r_n^B}{g_n}, \qquad r_n = \frac{\theta_n r_n^B + \lambda_n r_n^C}{g_n}. \tag{A40}$$

First suppose

$$\operatorname{cov}(r_n^B, r_m^B) = 0 \quad \text{and} \quad \operatorname{cov}(r_n^C, r_m^C) = 0. \tag{A41}$$

Using (A40) and Structural Orthogonality,

$$0 = \operatorname{cov}(r_n^B, \Delta Q_m) = \frac{\operatorname{cov}(r_n^B, r_m^C) - \operatorname{cov}(r_n^B, r_m^B)}{g_m}. \tag{A42}$$

Since $\operatorname{cov}(r_n^B, r_m^B) = 0$, this implies

$$\operatorname{cov}(r_n^B, r_m^C) = 0. \tag{A43}$$

Similarly,

$$0 = \operatorname{cov}(r_n^C, \Delta Q_m) = \frac{\operatorname{cov}(r_n^C, r_m^C) - \operatorname{cov}(r_n^C, r_m^B)}{g_m}, \tag{A44}$$

so

$$\operatorname{cov}(r_n^C, r_m^B) = 0. \tag{A45}$$

Then

$$\text{cov}(\Delta Q_n, \Delta Q_m) = \frac{\text{cov}(r_n^C, r_m^C) - \text{cov}(r_n^C, r_m^B) - \text{cov}(r_n^B, r_m^C) + \text{cov}(r_n^B, r_m^B)}{g_n g_m} = 0. \tag{A46}$$

Finally, using $r_n = r_n^B + \lambda_n \Delta Q_n$,

$$\text{cov}(r_n, r_m) = \text{cov}(r_n^B, r_m^B) + \lambda_m \text{cov}(r_n^B, \Delta Q_m) + \lambda_n \text{cov}(\Delta Q_n, r_m^B) + \lambda_n \lambda_m \text{cov}(\Delta Q_n, \Delta Q_m) = 0. \tag{A47}$$

The first term is zero by (A41), the middle terms are zero by Structural Orthogonality, and the last term is zero by (A46). Hence no common valuation shocks imply no common equilibrium return-flow variation.

Conversely, suppose

$$\text{cov}(r_n, r_m) = 0 \quad \text{and} \quad \text{cov}(\Delta Q_n, \Delta Q_m) = 0. \tag{A48}$$

From $r_n^B = r_n - \lambda_n \Delta Q_n$ and Structural Orthogonality,

$$0 = \text{cov}(r_n^B, \Delta Q_m) = \text{cov}(r_n, \Delta Q_m) - \lambda_n \text{cov}(\Delta Q_n, \Delta Q_m). \tag{A49}$$

Since $\text{cov}(\Delta Q_n, \Delta Q_m) = 0$, this implies

$$\text{cov}(r_n, \Delta Q_m) = 0. \tag{A50}$$

Interchanging n and m gives

$$\text{cov}(\Delta Q_n, r_m) = 0. \tag{A51}$$

Therefore,

$$\text{cov}(r_n^B, r_m^B) = \text{cov}(r_n - \lambda_n \Delta Q_n, r_m - \lambda_m \Delta Q_m) = 0. \tag{A52}$$

Similarly, from $r_n^C = r_n + \theta_n \Delta Q_n$,

$$\text{cov}(r_n^C, r_m^C) = \text{cov}(r_n + \theta_n \Delta Q_n, r_m + \theta_m \Delta Q_m) = 0. \tag{A53}$$

Thus no common equilibrium return-flow variation implies no common valuation shocks. Combining the two directions proves the result.

A.4 Proof of Proposition 3

Let $\mathbf{r}_t = (r_{1,t}, \ldots, r_{N,t})^\top$ and $\Delta\mathbf{Q}_t = (\Delta Q_{1,t}, \ldots, \Delta Q_{N,t})^\top$. Write

$$\mathbf{\Sigma}_r := \text{var}(\mathbf{r}_t), \qquad \mathbf{\Sigma}_Q := \text{var}(\Delta\mathbf{Q}_t). \tag{A54}$$

Let $L = \text{rank}(\mathbf{\Sigma}_r)$. Take the spectral decomposition

$$\mathbf{\Sigma}_r = \mathbf{V}\mathbf{D}\mathbf{V}^\top, \tag{A55}$$

where $\mathbf{V}$ is $N \times L$ with $\mathbf{V}^\top\mathbf{V} = \mathbf{I}_L$, and $\mathbf{D}$ is $L \times L$, diagonal, and positive definite. Now define

$$\mathbf{M} := \mathbf{D}^{1/2}\mathbf{V}^\top\mathbf{\Sigma}_Q\mathbf{V}\mathbf{D}^{1/2}. \tag{A56}$$

Take the spectral decomposition of $\mathbf{M}$:

$$\mathbf{M}\mathbf{G} = \mathbf{G}\mathbf{\Pi}, \tag{A57}$$

where $\mathbf{G}$ contains the eigenvectors associated with the retained $K \leq L$ eigenvalues, $\mathbf{G}^\top\mathbf{G} = \mathbf{I}_K$, and $\mathbf{\Pi} = \text{diag}(\pi_1, \ldots, \pi_K)$ with $\pi_1 \geq \cdots \geq \pi_K > 0$. Define the non-dollar currency weights by

$$\mathbf{W} := \mathbf{V}\mathbf{D}^{-1/2}\mathbf{G}. \tag{A58}$$

For each factor k, set the U.S. dollar weight to $w_{0,k} := -\sum_{n=1}^{N} w_{n,k}$, so that the factor is a zero-investment long-short currency portfolio.

The factor return vector is

$$\mathbf{r}_t^{\text{factor}} := \mathbf{W}^\top\mathbf{r}_t. \tag{A59}$$

Using (A55) and (A58),

$$\text{var}(\mathbf{r}_t^{\text{factor}}) = \mathbf{W}^\top\mathbf{\Sigma}_r\mathbf{W} = \mathbf{G}^\top\mathbf{D}^{-1/2}\mathbf{V}^\top\mathbf{V}\mathbf{D}\mathbf{V}^\top\mathbf{V}\mathbf{D}^{-1/2}\mathbf{G} = \mathbf{I}_K. \tag{A60}$$

Thus factor returns are uncorrelated and normalized to unit variance.

Next construct factor flows. Let $\boldsymbol{\beta}_n$ denote the multivariate beta of currency n's return with respect to the factor returns:

$$\boldsymbol{\beta}_n := \text{var}(\mathbf{r}_t^{\text{factor}})^{-1} \text{cov}(\mathbf{r}_t^{\text{factor}}, r_{n,t}). \tag{A61}$$

Equivalently, if $\mathbf{B}$ is the $N \times K$ matrix whose n-th row is $\boldsymbol{\beta}_n^\top$, then

$$\mathbf{B} = \boldsymbol{\Sigma}_r \mathbf{W}(\mathbf{W}^\top \boldsymbol{\Sigma}_r \mathbf{W})^{-1}. \tag{A62}$$

Since $\mathbf{W}^\top \boldsymbol{\Sigma}_r \mathbf{W} = \mathbf{I}_K$, this gives

$$\mathbf{B} = \boldsymbol{\Sigma}_r \mathbf{W} = \mathbf{V}\mathbf{D}\mathbf{V}^\top \mathbf{V}\mathbf{D}^{-1/2}\mathbf{G} = \mathbf{V}\mathbf{D}^{1/2}\mathbf{G}. \tag{A63}$$

The factor flow vector is the beta-weighted currency flow,

$$\Delta\mathbf{Q}_t^{\text{factor}} := \mathbf{B}^\top \Delta\mathbf{Q}_t = \mathbf{G}^\top \mathbf{D}^{1/2}\mathbf{V}^\top \Delta\mathbf{Q}_t. \tag{A64}$$

Therefore,

$$\text{var}(\Delta\mathbf{Q}_t^{\text{factor}}) = \mathbf{G}^\top \mathbf{D}^{1/2}\mathbf{V}^\top \boldsymbol{\Sigma}_Q \mathbf{V}\mathbf{D}^{1/2}\mathbf{G} = \boldsymbol{\Pi}. \tag{A65}$$

Since $\boldsymbol{\Pi}$ is diagonal, factor flows are uncorrelated.

Equations (A60) and (A65) prove (11) and (12). Moreover, under the normalization $\text{var}(r_{k,t}^{\text{factor}}) = 1$, the amount of trading-induced non-diversifiable risk captured by factor k is

$$\text{var}(r_{k,t}^{\text{factor}})\,\text{var}(\Delta Q_{k,t}^{\text{factor}}) = \pi_k.$$

Ordering eigenvalues π_k from largest to smallest therefore orders factors by (13). If a factor is rescaled, its return is rescaled and its beta-aggregated flow is rescaled inversely, so the product in (13) is unchanged.

Finally, uniqueness follows from the spectral decomposition of $\mathbf{M}$. When eigenvalues are distinct, the eigenvectors are unique up to sign; factor normalization is arbitrary; and when eigenvalues are repeated, any orthogonal rotation within the repeated-eigenvalue eigenspace yields an equivalent set of factors.

A.5 Proof of Proposition 4

It is enough to consider a change of numeraire from the U.S. dollar to currency N; any other numeraire is obtained by relabeling currencies. Let $\Delta\tilde{\mathbf{Q}}$ and $\tilde{\mathbf{r}}$ denote the currency-flow and return vectors under currency N as the numeraire. If ΔQ_n is the flow from USD into currency n, then the same trades expressed relative to currency N satisfy

$$\Delta\tilde{\mathbf{Q}} = \mathbf{C}\Delta\mathbf{Q}, \tag{A66}$$

where

$$\mathbf{C} := \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \\ -1 & -1 & \cdots & -1 & -1 \end{pmatrix}. \tag{A67}$$

The corresponding returns transform as

$$\tilde{\mathbf{r}} = \mathbf{C}^\top \mathbf{r}. \tag{A68}$$

The matrix $\mathbf{C}$ satisfies

$$\mathbf{C}\mathbf{C} = \mathbf{I}_N. \tag{A69}$$

This reflects the fact that changing the numeraire from USD to currency N, and then changing back, returns the original coordinates.

Let

$$\operatorname{var}(\mathbf{r}) = \mathbf{V}\mathbf{D}\mathbf{V}^\top \tag{A70}$$

be the spectral decomposition of the original return covariance matrix, where $\mathbf{V}^\top\mathbf{V} = \mathbf{I}_L$ and $\mathbf{D}$ is positive diagonal. Let

$$\operatorname{var}(\tilde{\mathbf{r}}) = \tilde{\mathbf{V}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^\top \tag{A71}$$

be the corresponding decomposition under the new numeraire. Since $\tilde{\mathbf{r}} = \mathbf{C}^\top\mathbf{r}$,

$$\operatorname{var}(\tilde{\mathbf{r}}) = \mathbf{C}^\top \operatorname{var}(\mathbf{r})\mathbf{C} = \mathbf{C}^\top\mathbf{V}\mathbf{D}\mathbf{V}^\top\mathbf{C}. \tag{A72}$$

Hence there exists an $L \times L$ orthogonal matrix $\mathbf{O}$, with $\mathbf{O}^\top \mathbf{O} = \mathbf{I}_L$, such that

$$\mathbf{C}^\top \mathbf{V} \mathbf{D}^{1/2} = \tilde{\mathbf{V}} \tilde{\mathbf{D}}^{1/2} \mathbf{O}. \tag{A73}$$

Under the original numeraire, the factor construction diagonalizes

$$\mathbf{M} := \mathbf{D}^{1/2} \mathbf{V}^\top \operatorname{var}(\Delta \mathbf{Q}) \mathbf{V} \mathbf{D}^{1/2}, \tag{A74}$$

so that

$$\mathbf{M}\mathbf{G} = \mathbf{G}\mathbf{\Pi}. \tag{A75}$$

Under the new numeraire, the analogous matrix is

$$\tilde{\mathbf{M}} := \tilde{\mathbf{D}}^{1/2} \tilde{\mathbf{V}}^\top \operatorname{var}(\Delta \tilde{\mathbf{Q}}) \tilde{\mathbf{V}} \tilde{\mathbf{D}}^{1/2}. \tag{A76}$$

Using (A66), (A69), and (A73),

$$\tilde{\mathbf{M}} = \mathbf{O} \mathbf{D}^{1/2} \mathbf{V}^\top \operatorname{var}(\Delta \mathbf{Q}) \mathbf{V} \mathbf{D}^{1/2} \mathbf{O}^\top = \mathbf{O} \mathbf{M} \mathbf{O}^\top. \tag{A77}$$

Therefore, if $\mathbf{M}\mathbf{G} = \mathbf{G}\mathbf{\Pi}$, then

$$\tilde{\mathbf{M}}(\mathbf{O}\mathbf{G}) = \mathbf{O}\mathbf{M}\mathbf{G} = (\mathbf{O}\mathbf{G})\mathbf{\Pi}. \tag{A78}$$

Thus we can take

$$\tilde{\mathbf{G}} = \mathbf{O}\mathbf{G}, \qquad \tilde{\mathbf{\Pi}} = \mathbf{\Pi}. \tag{A79}$$

The eigenvalues, and hence the ordering by trading-induced risk, are unchanged.

It remains to show that the constructed factor returns and flows are the same. Under the original numeraire, the factor-weight matrix is

$$\mathbf{W} = \mathbf{V} \mathbf{D}^{-1/2} \mathbf{G}. \tag{A80}$$

Under the new numeraire, it is

$$\tilde{\mathbf{W}} = \tilde{\mathbf{V}} \tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{G}}. \tag{A81}$$

Working with demeaned returns, which is what enters the covariance construction, write

$\mathbf{r} = \mathbf{Vf}$ for some L-dimensional random vector $\mathbf{f}$. From (A73),

$$\tilde{\mathbf{D}}^{-1/2}\tilde{\mathbf{V}}^\top \mathbf{C}^\top \mathbf{V} = \mathbf{O}\mathbf{D}^{-1/2}. \tag{A82}$$

Therefore,

$$\tilde{\mathbf{W}}^\top \tilde{\mathbf{r}} = \tilde{\mathbf{G}}^\top \tilde{\mathbf{D}}^{-1/2}\tilde{\mathbf{V}}^\top \mathbf{C}^\top \mathbf{r} = \mathbf{G}^\top \mathbf{O}^\top \tilde{\mathbf{D}}^{-1/2}\tilde{\mathbf{V}}^\top \mathbf{C}^\top \mathbf{V}\mathbf{f} = \mathbf{G}^\top \mathbf{D}^{-1/2}\mathbf{f} = \mathbf{W}^\top \mathbf{r}. \tag{A83}$$

Thus the factor returns are identical under either numeraire.

The same is true for factor flows. Under the original numeraire,

$$\Delta\mathbf{Q}^{\text{factor}} = \mathbf{G}^\top \mathbf{D}^{1/2}\mathbf{V}^\top \Delta\mathbf{Q}. \tag{A84}$$

Under the new numeraire, using $\tilde{\mathbf{G}} = \mathbf{OG}$, $\Delta\tilde{\mathbf{Q}} = \mathbf{C}\Delta\mathbf{Q}$, and $\tilde{\mathbf{D}}^{1/2}\tilde{\mathbf{V}}^\top = \mathbf{O}\mathbf{D}^{1/2}\mathbf{V}^\top \mathbf{C}$, we obtain

$$\Delta\tilde{\mathbf{Q}}^{\text{factor}} = \tilde{\mathbf{G}}^\top \tilde{\mathbf{D}}^{1/2}\tilde{\mathbf{V}}^\top \Delta\tilde{\mathbf{Q}} = \mathbf{G}^\top \mathbf{O}^\top \mathbf{O}\mathbf{D}^{1/2}\mathbf{V}^\top \mathbf{C}\mathbf{C}\Delta\mathbf{Q} = \mathbf{G}^\top \mathbf{D}^{1/2}\mathbf{V}^\top \Delta\mathbf{Q} = \Delta\mathbf{Q}^{\text{factor}}. \tag{A85}$$

Equations (A83) and (A85) show that changing the numeraire leaves the constructed factor returns and factor flows unchanged. Since $\tilde{\mathbf{\Pi}} = \mathbf{\Pi}$, the trading-induced risk shares are also unchanged. The only remaining indeterminacies are the usual sign and normalization choices, and rotations within eigenspaces associated with repeated eigenvalues.

A.6 Proof of Proposition 5

Suppress the time subscript. Let $\mathbf{B}$ denote the $N \times K$ matrix of currency factor exposures, with entry (n, k) equal to $\beta_{n,k}$. Let

$$\mathbf{\Lambda}^{\text{factor}} := \text{diag}(\lambda_1^{\text{factor}}, \dots, \lambda_K^{\text{factor}}).$$

By the beta aggregation of flows in (10), a currency-level customer-flow perturbation $d\Delta\mathbf{Q}$ induces factor-level customer flow

$$d\Delta\mathbf{Q}^{\text{factor}} = \mathbf{B}^\top d\Delta\mathbf{Q}. \tag{A86}$$

In the traded-risk-factor basis, the bank-side system (14) is diagonal, so holding factor-level bank valuation shocks fixed,

$$d\mathbf{r}^{\text{factor}} = \mathbf{\Lambda}^{\text{factor}} d\Delta\mathbf{Q}^{\text{factor}}. \tag{A87}$$

No-arbitrage maps factor repricing back to currency returns through the same factor exposures:

$$d\mathbf{r} = \mathbf{B}\, d\mathbf{r}^{\text{factor}}. \tag{A88}$$

Combining (A86), (A87), and (A88) gives

$$d\mathbf{r} = \mathbf{B}\mathbf{\Lambda}^{\text{factor}}\mathbf{B}^{\top} d\Delta\mathbf{Q}. \tag{A89}$$

By (1), the currency-level bank-side price sensitivity matrix satisfies

$$d\mathbf{r} = \mathbf{\Lambda}^{\text{ccy}} d\Delta\mathbf{Q}, \tag{A90}$$

holding bank-side valuation shocks fixed. Therefore,

$$\mathbf{\Lambda}^{\text{ccy}} = \mathbf{B}\mathbf{\Lambda}^{\text{factor}}\mathbf{B}^{\top}. \tag{A91}$$

Taking the (n, m) entry yields

$$\lambda_{n,m} = \sum_{k=1}^{K} \beta_{n,k}\lambda_k^{\text{factor}}\beta_{m,k}. \tag{A92}$$

Equivalently,

$$\lambda_{n,m} = \sum_{k=1}^{K} \frac{\partial \Delta Q_k^{\text{factor}}}{\partial \Delta Q_m} \frac{\partial r_k^{\text{factor}}}{\partial \Delta Q_k^{\text{factor}}} \frac{\partial r_n}{\partial r_k^{\text{factor}}} = \sum_{k=1}^{K} \beta_{m,k}\lambda_k^{\text{factor}}\beta_{n,k}. \tag{A93}$$

This proves (16).

B FX Derivatives and Spot-Equivalent Customer Flows

Foreign exchange trades can be executed in the spot market and in the derivatives market of forwards and swaps. Trading in FX derivatives creates spot-equivalent currency exposure for banks. Consider a customer-initiated trade in which the customer sells \$100 worth of JPY one-month forward against USD. In the absence of other trades, a bank that maintains a net neutral FX position and serves as the counterparty to this trade must satisfy the obligation to deliver \$100 in one month by setting aside $\$100/(1 + r^{\$}_{1M})$ today, where $r^{\$}_{1M}$ is the one-month USD risk-free rate. Similarly, the bank will sell $100/(1 + r^{JPY}_{1M})$ of JPY today to fund its USD purchase and to ensure FX neutrality when it receives the promised delivery from the customer. From the bank's perspective, a forward contract is therefore economically similar to a spot transaction, except that the implied spot exposure is slightly smaller than the notional amount because of discounting.

Because our main analysis measures the total customer FX flow absorbed by banks, we include trading flows in both the spot and non-spot derivatives markets.[17] In this appendix, we show that FX derivatives trading is both economically large and distinct from spot trading, making its inclusion important for measuring customer demand and banks' risk absorption.

We first examine gross trading activities. The Bank for International Settlements (BIS) has conducted the Triennial Central Bank Survey of Foreign Exchange and Over-the-Counter Derivatives since 2002. These data show that, globally, FX forward and swap activity involving USD is nearly twice as large as spot activity; see also Appendix Figure A1. The sheer size of derivative markets implies that excluding forwards and swaps would omit a substantial portion of customer FX demand.

Next, we examine spot-equivalent flows from FX derivatives using the CLS data. Specifically, we estimate the spot-equivalent customer flows implied by FX derivative transactions, net these flows within each currency-day, and compare the resulting flows with direct spot-market flows. Appendix Table A1 reports the correlation between net spot flows and net spot-equivalent derivative flows for the six major currencies in our sample. Correlations range from -0.62 to 0.17, indicating that derivative trading contains substantial variation

[17] We treat swaps as a spot transaction plus a forward contract. As discussed in Section 3.1, our results are effectively unchanged if swaps are excluded.

Figure A1: **FX Daily Turnover Against USD**

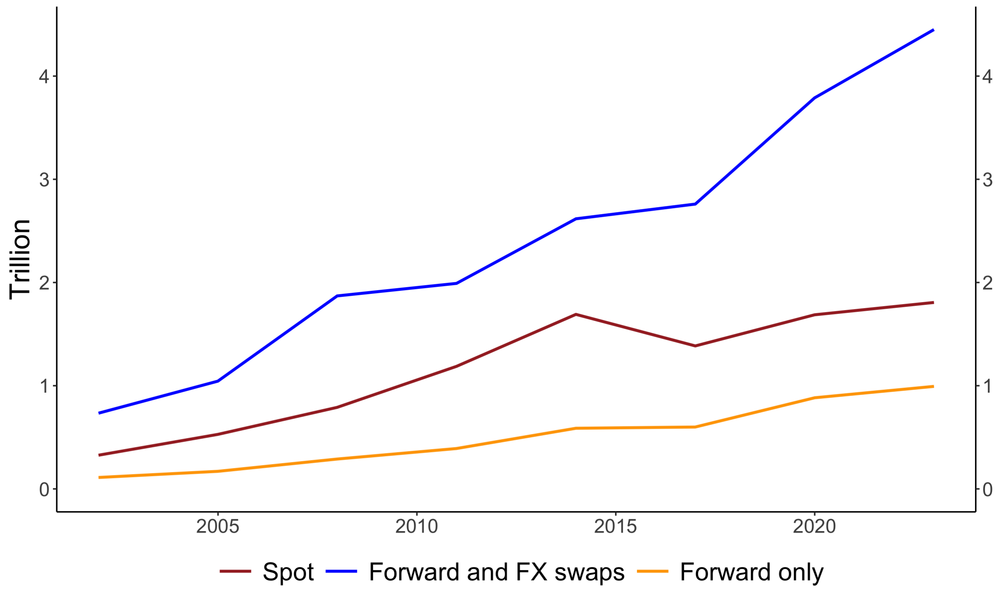

Notes: This figure plots the global daily volume of foreign exchange spot, forward, and FX swap transactions involving USD. Daily volume is calculated as the average of all trading days in April of the survey year. The survey is conducted triennially from 2001 to 2022 by the BIS.

Table A1: **Currency-Specific Correlation between Net Trading Flow in Spot vs. Non-Spot Derivatives**

AUD	CAD	CHF	EUR	GBP	JPY
-0.48	0.17	-0.54	-0.39	-0.62	-0.35

Notes: This table reports the correlation between net flows into individual currencies in the spot market and in the non-spot derivatives market.

not captured by spot transactions alone. Consequently, restricting attention to spot-market flows would provide an incomplete measure of customer demand and bank's risk exposure in FX markets.

C Implications of Reversal in Demand-Induced Price Response

This appendix quantifies how much of each factor's return variance can be attributed to the temporary price response of customer flows.

We start with the variance contribution. Customer flows may contain both information and demand components. The IV estimate $\lambda_k^{\text{factor}}$ isolates the demand-induced price response of factor k. Figure 2 shows that this response reverts over roughly one month. We use this reversion pattern to compute, across horizons, the share of factor return variance attributable to the temporary price response of customer flows.

Let $\sigma_{Q,k}$ denote the annualized volatility of customer flows into factor k, measured in billions of U.S. dollars, and let $\sigma_{r,k}$ denote the annualized volatility of factor returns. During an infinitesimal interval dt, the flow shock has standard deviation $\sigma_{Q,k}\sqrt{dt}$. Since $\lambda_k^{\text{factor}}$ measures the price response to one unit of flow, the contemporaneous price change generated by this flow shock has standard deviation $\lambda_k^{\text{factor}}\sigma_{Q,k}\sqrt{dt}$.

We approximate the price response as decaying linearly to zero within one month, $\tau = 1/12$ years. Thus, a flow shock arriving at time t still affects the price at terminal date T by the factor $\max\{0, 1-(T-t)/\tau\}$. Under this approximation, the share of cumulative return variance over horizon T attributable to the temporary price response of flows is

$$\frac{\int_0^T \left(\max\{0, 1-(T-t)/\tau\}\lambda_k^{\text{factor}}\sigma_{Q,k}\right)^2 dt}{\sigma_{r,k}^2 T}. \tag{A94}$$

With $\tau = 1/12$, this expression simplifies to

$$\frac{(\lambda_k^{\text{factor}})^2\sigma_{Q,k}^2}{\sigma_{r,k}^2}\begin{cases} 1-12T+48T^2, & 0 < T \leq \dfrac{1}{12}, \\ \dfrac{1}{36T}, & T > \dfrac{1}{12}. \end{cases} \tag{A95}$$

Figure A2 shows that for the three factors, flows explain roughly 10–25% of return variance at the 1-week horizon, 5–10% at one month, and rapidly less at longer horizons.

D Additional Tables

Figure A2: **Factor Flow Contribution to Return Variance**

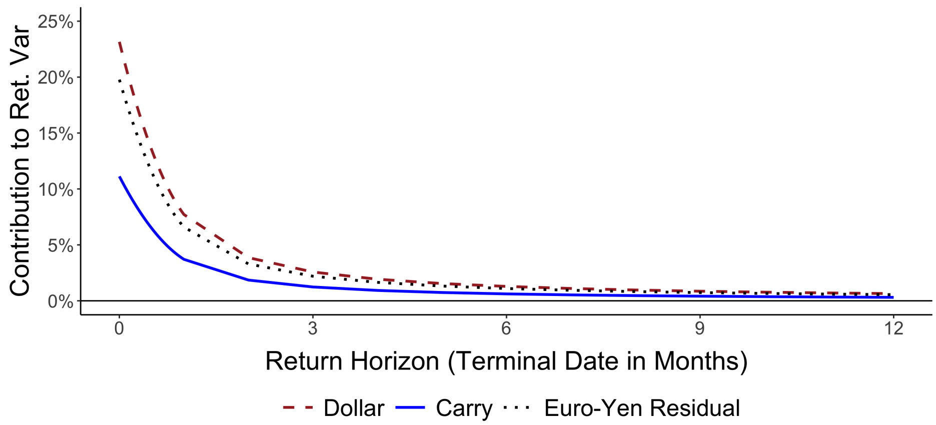

Notes: This figure displays the share of factor return variance attributable to the temporary price response of customer flows, computed using equation (A95). The calculation assumes that the contemporaneous price response of a flow shock decays linearly to zero within one month.

Table A2: **Correlation Between Traded Risk Factors in Full Sample vs. Subsamples**

		Factor 1	Factor 2	Factor 3
Return	Pre 2020	0.97	0.83	0.83
	Post 2020	1.00	0.97	0.89
Flow	Pre 2020	0.98	0.82	0.81
	Post 2020	0.99	0.96	0.81

Notes: In this table, we report the correlation between returns and flows of the traded risk factors constructed based on the full sample versus returns and flows of the traded risk factors constructed based on different subsamples. Pre-2020 refers to the sample period from September 2012 to December 2019, while post-2020 refers to the sample period from January 2020 to December 2023.