Demand Propagation Through Traded Risk Factors[*]

Yu An[†]   Amy W. Huber[‡]

First version: April 2024.
This version: July 2025. Most recent version here.

Abstract

We develop a risk-driven framework to quantify how demand shocks to one financial asset reprice others. Applying it to the foreign-exchange (FX) market, we find that intermediaries' aversion to absorbing non-diversifiable risk is a key channel for shock transmission. Using a novel technique that jointly analyzes trading flows and returns, we first identify three traded risk factors—Dollar, Carry, and a new factor linked to active euro–yen trading—that together explain 90% of the non-diversifiable risk borne by FX intermediaries. We then estimate their price sensitivity to demand using an instrumental-variable strategy enabled by our factor construction, which isolates factor-specific variation from correlated cross-section. We find that factor prices rise by 5–30 basis points per \$1 billion of net demand, highlighting intermediaries' limited risk-bearing capacity. This friction drives the price impact of demand shocks, first repricing non-diversifiable risks and then all currencies exposed to these risks. Finally, we trace how demand shocks ripple through 17 currencies and 6 asset classes, showing that interventions targeting a single currency can propagate globally, moving both other exchange rates and non-currency assets with currency risk exposure.

Keywords: Demand Propagation, Traded Risk Factors, Diversification, Intermediary, FX

1 Introduction

A defining feature of modern financial markets is their tight interlinkages. Episodes such as the Global Financial Crisis of 2007–2009 serve as stark reminders that shocks originating in one corner of the market can rapidly propagate across the globe (e.g., Allen and Gale, 2000; Pavlova and Rigobon, 2008). One class of shocks that arises frequently in asset markets are demand shocks. Although they do not reflect new fundamental information, demand shocks nonetheless move prices powerfully (e.g., Lee, Shleifer, and Thaler, 1991; Froot and Ramadorai, 2008; Koijen and Yogo, 2019). In an interconnected system, these shocks rarely affect only a single asset. For example, when a central bank intervenes in the foreign exchange market, the target currency moves—but so do many others. Which currencies react, and by how much? Answering these questions is essential for understanding how shocks propagate and for designing effective policy interventions and responses.

In this paper, we propose a risk-driven framework for quantifying demand propagation, and find that intermediaries' aversion to absorbing non-diversifiable risk is a key channel for shock transmission. We argue that when a single asset is hit by a demand shock, marginal investors—often intermediaries—end up carrying additional exposure in a small set of traded, non-diversifiable risk factors; intermediaries' limited risk-taking capacity means that shifts in factor demand result in factor prices change, which in turn reprice all assets that load on those factors.

The focus on non-diversifiable risks is motivated by the tight connections across financial assets. In principle, every asset could be affected by shocks to every other, thus flexibly capturing propagation across a market with N assets would require estimating N^2 cross-impact coefficients. However, genuine asset-specific demand variation is scarce: investor flows often arrive in correlated baskets, rebalancing typically affects multiple assets at once, and index-linked trades tend to move many markets together. Progress therefore hinges on imposing structure. Our structure derives from two well-established insights. First, asset prices co-move because investors care about non-diversifiable risks, which can be represented as factor portfolios or "factors" (Markowitz, 1952; Ross, 1976; Kozak, Nagel, and Santosh, 2018). Second, these factors are inelastic: small quantity changes generate large price movements (Gabaix and Koijen, 2021; Li and Lin, 2022). Our contribution is to empirically link factor inelasticities with each asset's exposure to those factors, and to identify intermediaries' limited risk-bearing capacity as a key friction in the risk-consistent transmission of shocks.

We apply our framework to study demand propagation in foreign exchange (FX). Demand shocks, such as those caused by central-bank interventions, corporate hedging, and index rebalancing, hit currencies almost daily. At the same time, no currency market operates in

isolation. Sophisticated intermediaries absorb customer trades and pass any resulting inventory risk among themselves through a dense inter-dealer network. This constant reshuffling leaves returns to align with a handful of common factors (Lustig, Roussanov, and Verdelhan, 2011), underscoring the role of non-diversifiable risks in driving currency co-movement. What has been missing is data on the underlying trading flows. A novel dataset that records the net positions handled by more than 70 major intermediaries between 17 currencies fills this gap, which we harness to study the risks FX intermediaries truly bear.

Our inquiry starts with identifying the risk factors actually transmit demand shocks. Risk factors are typically proposed to explain common variation in returns (Ross, 1976). But a singular focus on returns ignores quantities, and the resulting risk factors need not correspond to those that investors actually trade. We therefore develop a new method that jointly analyzes the weekly trading flows between customers and intermediaries across 17 major currencies and the corresponding panel of currency returns, uncovering what we call "*traded risk factors*": portfolios that explain the majority of non-diversifiable risk borne by intermediaries when absorbing trading flows, and that therefore serve as the principal conduits for demand-shock propagation.[1]

Our method reveals that the top three traded FX factors account for 90% of the trading-induced non-diversifiable risks in our sample, are economically interpretable, and remain stable over time. The two most traded FX factors resemble the well-known Dollar and Carry factors, yet our method recovers the Carry factor without first sorting currencies by interest rates, as in Lustig, Roussanov, and Verdelhan (2011). It also exposes intermediaries' hidden risk exposures: from 2012 to 2023 they accumulated about 0.8 trillion of Carry exposure. The third factor, the Euro-Yen Residual, captures the risk intermediaries bear when accommodating active customer trading between the euro area and Japan, after hedging Dollar and Carry exposures. This new factor delivers a Sharpe ratio comparable to that of the Carry factor. Finally, unlike the standard "flow-only" method (e.g., Lo and Wang, 2000; Hasbrouck and Seppi, 2001) applied to FX, our approach combines flows and returns and thus extracts genuine risk factors that customers trade, not just portfolios of actively traded currency pairs.

Next, we investigate how the price of each traded FX factor responds to its own trading-induced risks. To estimate this price sensitivity from observed trading flows and returns, we use instrumental variables (IVs). A valid instrument must satisfy three conditions: rel-

---

[1]Our notion of "traded" risk factors differs from the conventional idea of "tradable" factors. Tradable factors correspond to risks that are replicable by portfolios that investors could buy or sell. Non-tradable factors, by contrast, reflect risks that are observable but for which no direct trading vehicle exists—global macroeconomic conditions are a typical example. Crucially, labeling a factor "tradable" does not imply that the associated replicating portfolio is, in fact, actively traded in equilibrium.

evance, exogeneity, and exclusion restriction. In our context, the ideal instrument should affect trading demand without carrying information about fundamentals, and crucially, the instrument should affect a factor's price only though demand for that factor. This last requirement is difficult to satisfy in general. As financial assets are tightly linked, any instrument that shifts demand for one asset is likely to shift demand for others that are not included in the regression. Such spill-overs of demands feed back into the original asset's price via cross-impacts, contaminating the estimate of own price sensitivity. To address this issue, we orthogonalize our traded risk factors so that *both* factor returns and factor trading flows are mutually uncorrelated. Put differently, each factor is an independent source of risk that investors also trade independently in equilibrium. This joint orthogonality ensures that each factor behave like an independent asset: a demand shock to one leaves the prices of the rest unchanged. Hence, a consistent estimate of each factor's price sensitivity follows simply from regressing that factor's return on its own instrumented demand.[2]

We use as instrumental variables the *week-ahead announcements* of the offering amount at upcoming sovereign bond auctions in the U.S., Australia, Canada, France, Germany, Italy, Japan, and the U.K. These sovereign auctions often attract foreign investors who need to convert currencies to participate, making the instruments relevant. Importantly, because these auctions are typically forward-guided, the week-ahead announcements contain limited new information, making the instruments plausibly exogenous and satisfying the exclusion restriction. For instance, in the U.S., the Treasury Borrowing Advisory Committee (TBAC) releases two-quarter-ahead recommendation on auction amounts; the subsequent week-ahead announcements and the eventual auctions exhibit little deviation from these recommendations (Rigon, 2024). Although the timing and the amount of bonds auctioned are well anticipated, the associated demand shocks could still affect prices when realized (Vayanos, 2021; Hartzmark and Solomon, 2024). We find that compensating intermediaries for absorbing a \$1 billion non-diversifiable demand shock requires price increases of 5 basis points (bps) for the Dollar, 9 bps for the Carry, and 29 bps for the Euro-Yen Residual. These contemporaneous price responses fully revert within a month, supporting the interpretation that our instruments capture uninformed shocks. As the price impact from trading is mean-reverting, trading volatility explains a sizable share of short-term return variance—about 10–35% of the 1-week return, but considerably less at longer horizons—5–15% of the 1-month return.

---

[2]Fuchs, Fukuda, and Neuhann (2025) and Haddad, He, Huebner, Kondor, and Loualiche (2025) both point out that cross-asset linkages can complicate IV estimation. Haddad, He, Huebner, Kondor, and Loualiche (2025) tackle the problem by regressing each factor's return on the instrumented demands for all factors simultaneously. Our orthogonal factor design achieves the same identification without first specifying the entire universe of risk factors.

Our estimated factor price sensitivity directly captures the intermediaries' aversion to absorbing additional risk in that factor, and our results imply that FX intermediaries have relatively limited risk-bearing capacity. Such intermediation friction gives rise to demand shocks' price impact. In fact, the traded FX factors exhibit higher price sensitivity to demand shocks than the U.S. equities market factor.[3] Supporting the interpretation that higher price sensitivity reflects intermediaries' more limited risk-bearing capacity, we also find evidence of state-dependence. Specifically, when intermediaries' public equity returns are high, the Dollar factor's price sensitivity is low. As equity returns likely capture variations in intermediaries' wealth, our findings are consistent with greater wealth leading to greater willingness to bear risk (He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014). One potential reason for FX intermediaries' limited risk-bearing capacity is the relatively limited arbitrage capital in this highly specialized market. Although there is large daily trading volume in FX, up to 75% of trades are between intermediaries (BIS, 2022), leaving a more modest pool of capital to absorb demand shocks. Differences in arbitrage capital may also explain the cross-factor variation in price sensitivity, with lesser-known factors such as the Euro-Yen Residual attracting less arbitrage capital and exhibiting greater price sensitivity.

Having identified the most traded FX factors and each factor's price sensitivity to risk, we use these findings to trace out demand propagation across currencies. When intermediaries accommodate a demand shock to one currency, they bear additional non-diversifiable risks captured by the traded FX factors. These risks then affect the prices of the traded FX factors in proportion to these factors' price sensitivities. Finally, the law of one price ensures that prices of other currencies that share those exposures also adjust and such exposures can be measured as betas or factor loadings. We quantify this demand propagation through cross-multipliers, which measure how a demand shock to one currency affects the price of another, holding the demand shocks to all other currencies constant.[4]

We uncover rich patterns of cross-currency substitution arising from exposures to the three traded risk factors. Currencies exhibit strong demand propagation when they share the same sign of loading to a factor and modest propagation when they have opposite signs. For instance, we find a large cross-multiplier between the Australian dollar (AUD) and

---

[3]Gabaix and Koijen (2021) find that a 1% larger trading demand shock to the entire U.S. stock market increases price by 5%. Such a shock can be interpreted as a shock to the market factor. The average market capitalization between 2012 and 2022 is about $31.7 trillion. A $1 billion demand shock in our sample period therefore raises the price of the market factor by about 2 bps.

[4]Because *idiosyncratic* risk does not transmit demand shock across different currencies, our three factors home in on the part that does: *systematic* risk. That said, our results should not be interpreted as capturing the total multipliers across all currencies. Doing so would require measuring both price responses to systematic risk *and* price responses to idiosyncratic risk while holding systematic risk constant.

the Canadian dollar (CAD) because both currencies have the same sign of loadings on all three traded FX factors, making them close "substitutes" whose prices co-move strongly. In contrast, the cross-multiplier between the Japanese yen (JPY) and either AUD or CAD is small because JPY has the opposite loading on the Carry factor, allowing these currencies to hedge each other by reducing intermediaries' exposure to the Carry factor. Similarly, while the euro (EUR) and JPY are both low-interest-rate currencies and act as "substitutes" with respect to the Carry factor, they are on opposite sides of the Euro-Yen Residual factor, making them "complements" for that factor. As a result, we estimate only a modest cross-multiplier, implying muted demand propagation between EUR and JPY.

Additionally, we find that five major non-currency assets also load on the traded FX factors, allowing demand shocks to propagate across markets through shared currency risk. We study U.S. Treasury bonds (Treasurys), corporate bonds, options, CDS, and commodities. The traded FX factors explain approximately 30% of each asset class' return variance. Consequently, a demand shock to, say, corporate bonds generates non-diversifiable risks as captured by the traded FX factors. These risks affect the prices of traded FX factors and, in turn, the prices of options and other assets that load on the affected factors. Demand shocks transmitted through the traded FX factors have the smallest price effect on U.S. Treasurys, consistent with the depth and liquidity of the Treasurys market. Notably, only Treasurys exhibit negative cross-multipliers with other assets, reflecting Treasurys' "safe haven" status. In our framework, this safe-haven property arises because Treasurys load on the Carry factor with a uniquely negative sign, making them an effective hedge against other asset classes during shifts between "risk-on" and "risk-off" regimes.

Our paper contributes to the understanding of demand propagation across assets marked by high degree of interconnection, substitutability, and complementarity. In essence, we introduce a new structure to the $N \times N$ cross-impact coefficients by mapping them onto a small set of orthogonal risk factors that govern price co-movement. This risk-based structure follows naturally from the canonical asset-pricing insight that risk drives returns (Markowitz, 1952; Ross, 1976). In doing so, our approach complements and contrasts with other structural approaches of studying demand shocks. One existing approach maps the effect of demand to asset characteristics via micro-founded demand systems (e.g., Koijen and Yogo, 2019, 2020; Bretscher, Schmid, Sen, and Sharma, 2022; Chaudhary, Fu, and Li, 2023; Jiang, Richmond, and Zhang, 2024; Haddad, He, Huebner, Kondor, and Loualiche, 2025). Another approach links the price effect of demand shocks to pairwise return covariances (e.g., Vayanos and Vila, 2021; Kodres and Pritsker, 2002; Pasquariello and Vega, 2015; Davis, Kargar, and Li, 2023; Greenwood, Hanson, and Vayanos, 2023; Jansen, Li, and Schmid, 2024). Similar to these approaches, we take seriously the factor structure in asset returns, which is shown to

be empirically important even in the presence of noise trading (Kozak, Nagel, and Santosh, 2018). At the same time, we innovate on two key dimensions. First, because we are interested in propagation across assets, we directly study the effect of demand shocks to traded, non-diversifiable risks. Second, we allow possibly different risk-bearing capacity, and thereby price sensitivity, toward different risks. Consequently, we are able to generate flexible cross-asset dynamics that allow assets to be substitutes with respect to one risk but complements with respect to another.

Our paper also extends the literature on exchange rates by developing a novel approach to quantifying the price response of trading flows through risks. Beyond conveying information (e.g., Evans and Lyons, 2002; Pasquariello, 2007; Froot and Ramadorai, 2008), trading influences prices by increasing the non-diversifiable risks that marginal investors must bear. Our contribution is to recover risk factors that investors empirically deem important by analyzing their trading behavior together with return data. This revealed-preference approach differs from, and complements, the literature's typical method of conjecturing relevant state variables based on economic intuition, constructing factors from those variables, and then testing these factors' cross-sectional pricing power.[5] We find that the two most traded FX factors, the Dollar and the Carry, are similar to what price unconditional FX returns (Lustig, Roussanov, and Verdelhan, 2011). We moreover introduce a new Euro-Yen Residual factor, which delivers a Sharpe ratio comparable to that of the Carry factor and is priced conditionally on demand shocks. Finally, we uncover new evidence demonstrating the pivotal role of intermediation friction in how risks are priced in FX (e.g., Gabaix and Maggiori, 2015, Itskhoki and Mukhin, 2021). Our findings complement existing research on priced risk factors in FX markets (e.g., Bansal and Dahlquist, 2000; Lustig and Verdelhan, 2007; Hassan and Mano, 2018; Korsaye, Trojani, and Vedolin, 2023) and offer fresh insights into the role of trading-induced risks in driving price co-movements across currencies, between FX and other asset markets, as well as in transmitting 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; Liao and Zhang, 2025).

More generally, our paper augments the intermediary asset pricing and the microstructure literature, both of which emphasize intermediaries' limited risk-bearing or balance-sheet capacity as a driver of asset price responses to customers' demand shocks (e.g., Ho and Stoll, 1981; Grossman and Miller, 1988; Gabaix and Maggiori, 2015; He and Krishnamurthy, 2017;

---

[5]For example, Fama and French (1993) identify size and value as key state variables for determining 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.

Kondor and Vayanos, 2019; Haddad and Muir, 2021; Du, Hébert, and Huber, 2023; Du, Hébert, and Li, 2023). While we share this focus on intermediaries and the frictions they face, our approach differs in emphasizing that intermediaries' pricing decisions are shaped by non-diversifiable risks aggregated across all assets, rather than analyzing the risks of individual assets in isolation. In this sense, our perspective aligns with the foundational insights of Markowitz (1952), Sharpe (1964), and Lintner (1965), where non-diversifiable risks are the primary concern in asset price determination.

The next section presents our theoretical framework. Section 3 introduces the data sources and Section 4 identifies the traded FX factors. Section 5 examines the unconditional and conditional pricing properties of the traded FX factors. Section 6 explores how these factors propagate demand shocks across currencies and other asset classes. Section 7 concludes.

2 Theoretical Framework

This section shows when the complex problem of currency-level demand propagation can be reduced to tracking a few traded risk factors. We begin by revisiting the standard equilibrium model and show how non-diversifiable risks lead demand shocks to affect prices of both individual currencies and arbitrary currency portfolios (factors). We then introduce two improvements. (i) Motivated by empirical evidence that portfolios exhibit heterogeneous price sensitivities to demand shocks, we relax the standard model's assumption of a single, common price-sensitivity parameter and instead allow each factor to have its own sensitivity. (ii) We construct these factors to maximally explain the non-diversifiable risk intermediaries bear when absorbing trading flow, while requiring their equilibrium returns and flows to be mutually uncorrelated; each factor thus becomes an independent channel for cross-currency demand propagation. Finally, we show that these factor-level sensitivities, combined with the absence of arbitrage, yield clear analytical formula for how demand shocks propagate across currencies.

2.1 Standard Equilibrium Model

Consider a two-period setting, $t = 0$ and $t = 1$, with $N + 1$ currencies. The last currency serves as the numeraire, so all payoffs are denominated in its units. Appendix A.3 shows that our factor construction is invariant to the numeraire choice; hence, without loss of generality, we take the U.S. dollar (USD) as the numeraire throughout. At $t = 0$ customers can buy or sell any pair of currencies. A continuum of competitive intermediaries with mass μ absorbs all customer orders and clears the market. Each intermediary has constant absolute risk

aversion (CARA) utility with coefficient γ. For each non-USD currency $n = 1, \ldots, N$, let r_n denote its excess return from $t = 0$ to $t = 1$. Formally, r_n is the payoff from borrowing one unit of the USD at the risk-free rate, converting the proceeds into currency n at $t = 0$, investing at currency n's risk-free rate until $t = 1$, and then converting the position back into the USD at $t = 1$.

We study customer demand shocks that arrive at $t = 0$. These shocks are uninformed: they are independent of currency payoffs that materialize at $t = 1$. Because intermediaries have limited risk-bearing capacity, the price of currency n at time 0 moves from P_n to $P_n (1 + \Delta p_n)$, where Δp_n is the percentage price change—or price impact—on currency n. Customers may trade any pair of currencies. For accounting purposes we rewrite every transaction as a flow into or out of each non-USD currency against the USD.[6] Let ΔQ_n denote the resulting net purchase of currency n (measured in units of the USD and valued at the pre-shock price P_n).

The equilibrium prices at $t = 0$ must adjust so that competitive intermediaries are willing to absorb the net customer demand. Let each intermediary take a position y_n (measured in units of the USD) in currency n, which in equilibrium equals $-\Delta Q_n/\mu$. Holding one additional USD's worth of currency n exposes the intermediary to an extra payoff r_n at $t = 1$, while the purchase price at time 0 has already risen by the percentage change Δp_n. Compounding that price change forward to time 1 at the USD's gross risk-free rate R_F gives the intermediary's objective

$$\{-\Delta Q_1/\mu, \ldots, -\Delta Q_N/\mu\} = \arg \max_{\{y_1, \ldots, y_N\}} \mathbb{E}\left[-\gamma \exp\left(-\sum_{n=1}^N y_n (r_n - R_F \Delta p_n)\right)\right]. \quad (1)$$

The first-order condition to (1) implies that the price impact of currency n satisfies[7]

$$\Delta p_n = \lambda \left[\text{cov}(r_n, r_1)\Delta Q_1 + \text{cov}(r_n, r_2)\Delta Q_2 + \cdots + \text{cov}(r_n, r_N)\Delta Q_N\right], \quad (2)$$

where $\lambda := \gamma/(\mu R_F)$ approximates the per-capita risk aversion of intermediaries. The division by the gross risk-free rate R_F appears only because the model is discrete; in continuous time it disappears. Equation (2) shows that the price impact of asset n is a return-covariance-weighted linear combination of demand shocks across all currencies—the standard mean-variance benchmark for price impact (see Kozak, Nagel, and Santosh, 2018 and the

---

[6]If a customer buys currency n by selling currency m, we record $+\Delta Q_n$ for currency n and $-\Delta Q_m$ for currency m.

[7]Define the $N \times 1$ vectors $\Delta \mathbf{Q} = (\Delta Q_1, \ldots, \Delta Q_N)^\top$, $\mathbf{r} = (r_1, \ldots, r_N)^\top$, and $\Delta \mathbf{p} = (\Delta p_1, \ldots, \Delta p_N)^\top$. The first-order condition of (1) then takes the form $R_F \Delta \mathbf{p} - \mathbb{E}[\mathbf{r}] = \gamma \text{var}(\mathbf{r})\Delta \mathbf{Q}/\mu$. Note that $\Delta \mathbf{Q} = \mathbf{0}$ gives $\Delta \mathbf{p} = \mathbf{0}$. Taking the difference then gives $\Delta \mathbf{p} = (\gamma/(R_F \mu))\text{var}(\mathbf{r})\Delta \mathbf{Q}$, which elementwise corresponds to (2).

literature review by Rostek and Yoon, 2023).

Next, we show that the equilibrium price-impact formula in (2) extends from single currencies to any portfolio of currencies—i.e., to any factor. Fix a factor k with currency weights $w_{n,k}$ for $n = 1, \ldots, N$.[8] Let $\Delta p_k^{\text{factor}}$ and r_k^{factor} denote, respectively, the factor's price impact at $t = 0$ and its return from $t = 0$ to $t = 1$. Then

$$\Delta p_k^{\text{factor}} = w_{1,k} \Delta p_1 + \cdots + w_{N,k} \Delta p_N \tag{3}$$

$$= \lambda \left[\text{cov} \left(\sum_{n=1}^{N} w_{n,k} r_n, r_1 \right) \Delta Q_1 + \cdots + \text{cov} \left(\sum_{n=1}^{N} w_{n,k} r_n, r_N \right) \Delta Q_N \right] \tag{4}$$

$$= \lambda \left[\text{cov} \left(r_k^{\text{factor}}, r_1 \right) \Delta Q_1 + \cdots + \text{cov} \left(r_k^{\text{factor}}, r_N \right) \Delta Q_N \right]. \tag{5}$$

Equation (3) uses the Law Of One Price (LOOP) to express the factor's price impact as the weighted sum of its constituent currencies' impacts. Substituting the currency-level expression from (2) yields (4). Finally, recognizing that $r_k^{\text{factor}} = \sum_{n=1}^{N} w_{n,k} r_n$ and applying the LOOP once more gives (5).

2.2 Factor-Level Generalization

The standard equilibrium benchmark (2) above captures all $N \times N$ cross-currency impacts with one parameter, λ. We now relax this restriction by allowing a small set of K factor-specific price sensitivities. This generalization delivers the empirical flexibility needed to match the data while keeping the analysis firmly grounded in equilibrium theory.

Consider K risk factors whose returns are mutually uncorrelated, $\text{cov} \left(r_k^{\text{factor}}, r_j^{\text{factor}} \right) = 0$ for all $k \neq j$, which can be achieved by many orthogonalization techniques. Under the benchmark model (2), uncorrelated return alone eliminates cross-impacts. In our generalized setting this is insufficient, and we thus add an extra orthogonality condition in Section 2.4 to ensure factors remain cross-impact-free. For now, because factors have uncorrelated returns, the coefficient from a univariate regression of a currency's return on one factor equals the corresponding coefficient in a full multiple-factor regression. Thus the beta of currency n on factor k is simply $\beta_{n,k} = \text{var} \left(r_k^{\text{factor}} \right)^{-1} \text{cov} \left(r_k^{\text{factor}}, r_n \right)$. Substituting this expression into equation (5) yields

$$\Delta p_k^{\text{factor}} = \lambda \text{var} \left(r_k^{\text{factor}} \right) \left[\beta_{1,k} \Delta Q_1 + \cdots + \beta_{N,k} \Delta Q_N \right]. \tag{6}$$

Next, note that the linear combination $\beta_{1,k} \Delta Q_1 + \cdots + \beta_{N,k} \Delta Q_N$ represents the net customer demand for factor k. When intermediaries absorb a currency-level demand shock,

---

[8]Because every trade is financed in the USD, the implied portfolio weight on the USD is $-\sum_{n=1}^{N} w_{n,k}$.

ΔQ_n, their exposure to factor k rises by $\Delta Q_n \beta_{n,k}$, while the residual risk is orthogonal to that factor. Summing these beta-scaled exposures across all N currencies[9] therefore leaves a *non-diversifiable* demand shock of magnitude $\beta_{1,k}\Delta Q_1 + \cdots + \beta_{N,k}\Delta Q_N$.

This weighting scheme is exactly the portfolio-beta calculation used in risk management: each position is multiplied by its factor loading, and the results are added up to obtain the book's aggregate factor exposure. It is also identical to the cross-hedging ratio that standard texts derive for offsetting the risk in a single asset with an index future (see section 3.6 of Hull, 2022). Accordingly, define

$$\Delta Q_k^{\text{factor}} := \beta_{1,k}\Delta Q_1 + \cdots + \beta_{N,k}\Delta Q_N \tag{7}$$

as demand for factor k, and rewrite (6) as

$$\Delta p_k^{\text{factor}} = \lambda \text{var}\left(r_k^{\text{factor}}\right)\Delta Q_k^{\text{factor}}. \tag{8}$$

Rearranging equation (8) gives

$$\frac{\Delta p_k^{\text{factor}}}{\text{var}\left(r_k^{\text{factor}}\right)\Delta Q_k^{\text{factor}}} = \lambda = \frac{\gamma}{\mu R_F}. \tag{9}$$

Here, $\Delta p_k^{\text{factor}}$ represents the price impact of factor k at time 0. The denominator, $\text{var}\left(r_k^{\text{factor}}\right)\Delta Q_k^{\text{factor}}$ is the incremental quantity of risk created by the marginal demand shock. Consequently, the ratio measures the price sensitivity to risks induced by demand shocks and captures the *marginal* risk–return tradeoff conditional on trading. This concept extends the canonical price of risk, which reflects the *unconditional* risk–return tradeoff. Equation (9) further shows that the ratio equals intermediaries' per-capita risk aversion and is therefore invariant to portfolio leverage.[10]

We also note that equation (9) implies that for two different factors $k \neq j$, their price sensitivity to risk must be the same, i.e.,

$$\frac{\Delta p_k^{\text{factor}}}{\text{var}\left(r_k^{\text{factor}}\right)\Delta Q_k^{\text{factor}}} = \frac{\Delta p_j^{\text{factor}}}{\text{var}\left(r_j^{\text{factor}}\right)\Delta Q_j^{\text{factor}}}. \tag{10}$$

The restriction follows from mean–variance utility: a representative intermediary has the same marginal utility of wealth no matter which factor delivers that wealth. This property

---

[9]Our model uses a representative intermediary to accommodate all customer trades. In practice, such netting across currencies could also occur through interdealer trading.

[10]For instance, doubling every position—i.e., multiplying the weights $w_{n,k}$ by 2—doubles the factor's price impact $\Delta p_k^{\text{factor}}$, halves the factor demand $\Delta Q_k^{\text{factor}}$ (because each asset's beta to the levered factor is halved), and quadruples the return variance $\text{var}(r_k^{\text{factor}})$, leaving the ratio in (9) unchanged.

holds not only under the CARA specification we adopt, but more generally under CRRA or recursive preferences, as such utilities can generate time-varying risk aversion through state variables, but they preserve the cross-sectional equality implied by mean–variance pricing.

Because a single-lambda world is often too coarse, we relax (10) by allowing factor-specific sensitivities,

$$\frac{\Delta p_k^{\text{factor}}}{\text{var}\left(r_k^{\text{factor}}\right)\Delta Q_k^{\text{factor}}} = \lambda_k. \tag{11}$$

The motivation are twofold. Theoretically, modern asset-pricing research permits distinct unconditional risk–return trade-offs across factors (Fama and MacBeth, 1973). If unconditional trade-offs differ, it is natural to permit the conditional trade-offs in (11) to differ as well. Within the CARA world, $\lambda = \gamma/(\mu R_F)$, so heterogeneous trade-offs can arise if the pricing kernel varies by risk dimension (different γ) or if intermediaries are segmented across risks (different μ). Dynamic models can also produce distinct λ_k's endogenously; An and Zheng (2025) show this for CARA agents when trading demand is predictable.

Empirically, our ultimate goal is to recover the full $N \times N$ demand-propagation matrix. The standard benchmark in (2) collapses that matrix to a single degree of freedom, predicting the same price sensitivity whether the shock loads on aggregate market risk, a long–short style factor, or pure idiosyncratic risk. This stark implication conflicts with a large demand-based literature that finds markedly different price sensitivities for different portfolios (see Table 1 in Gabaix and Koijen, 2021 for a review).

Under our model (11), demand for one factor moves only its own price, not the prices of other factors. Our framework therefore contains only K sensitivity parameters—one λ_k per factor. By contrast, in a $K \times K$ specification such as Haddad, He, Huebner, Kondor, and Loualiche (2025), a shock to any factor can move the prices of all K factors, generating K^2 sensitivities. Both designs have empirical appeal. Our orthogonal, one-parameter-per-factor approach maximizes parsimony and assigns each λ_k a clear equilibrium interpretation: the marginal risk–return trade-off associated with factor k. It also permits unbiased estimation of λ_k one factor at a time, without committing to the full set of risk factors that might be relevant.

2.3 Demand Propagation Across Currencies

We now use the LOOP to connect factor-level price sensitivity with demand propagation across individual currencies. Consider the scenario where currency m experiences a \$1 demand shock (a one-dollar increase in ΔQ_m), while customers' demand for all other currencies remains constant. First, as in equation (7), this additional \$1 demand shock to currency m raises the demand shock $\Delta Q_k^{\text{factor}}$ to factor k by an amount $\beta_{m,k}$. Second, the change in

factor-k demand alters its price by $\Delta p_k^{\text{factor}} = \lambda_k \text{var}\left(r_k^{\text{factor}}\right)$, as in equation (11). Finally, changes in factor-k price $\Delta p_k^{\text{factor}}$ affect currency-n price Δp_n through the LOOP, with the sensitivity being $\beta_{n,k}$. The partial derivative $\partial \Delta p_n / \partial \Delta Q_m$ measures this demand propagation and is referred to as the "cross-multiplier." Proposition 1 derives the model-implied cross-multiplier, and Appendix A.1 provides a proof.

PROPOSITION 1 (Demand propagation). *The cross-multiplier between currencies n and m is:*

$$\frac{\partial \Delta p_n}{\partial \Delta Q_m} = \sum_{k=1}^{K} \frac{\partial \Delta Q_k^{factor}}{\partial \Delta Q_m} \times \frac{\partial \Delta p_k^{factor}}{\partial \Delta Q_k^{factor}} \times \frac{\partial \Delta p_n}{\partial \Delta p_k^{factor}} = \sum_{k=1}^{K} \beta_{m,k} \times \lambda_k \, \text{var}\left(r_k^{factor}\right) \times \beta_{n,k}. \quad (12)$$

The model-implied cross-multiplier has two features. First, the own-multiplier

$$\frac{\partial \Delta p_n}{\partial \Delta Q_n} = \sum_{k=1}^{K} \beta_{n,k}^2 \times \lambda_k \, \text{var}\left(r_k^{\text{factor}}\right) \quad (13)$$

is always positive so long as $\lambda_k > 0$. Positive λ_k indicates that intermediaries are averse to bearing trading-induced risks rather than risk-seeking. On the other hand, the cross-multiplier between two currencies can be negative when the currencies have opposite beta signs with respect to some factor—a form of complementarity. We return to this point empirically in Section 6.

Second, the cross-multiplier *as channeled through traded risk factors* is symmetric:

$$\frac{\partial \Delta p_n}{\partial \Delta Q_m} = \frac{\partial \Delta p_m}{\partial \Delta Q_n}, \quad (14)$$

so a \$1 demand shock in currency m moves currency n's price by the same percentage that a \$1 shock in currency n moves currency m's price. This symmetry arises because

$$\frac{\partial \Delta Q_k^{\text{factor}}}{\partial \Delta Q_n} = \beta_{n,k} = \frac{\partial \Delta p_n}{\partial \Delta p_k^{\text{factor}}}. \quad (15)$$

The first equality, relating currency to factors in terms of quantity, follows from our portfolio theory (7), while the second equality, relating currency to factors in terms of price, results from the law of one price. Both relationships are governed by the beta of currency n to factor k, which gives rise to the symmetry.

Under the standard equilibrium benchmark (2), the cross-impact between two currencies is

$$\frac{\partial \Delta p_n}{\partial \Delta Q_m} = \lambda \, \text{cov}(r_n, r_m), \quad (16)$$

12

which depends on a single parameter λ and the pairwise return covariance. Compared with our K-factor propagation in Proposition 1, the equal-λ restriction collapses the underlying factor structure of returns, leaving the cross-impact to depend solely on pairwise covariances.

2.4 Constructing Traded Risk Factors

The discussion thus far has focused on the asset-pricing implication of uninformed demand shocks. We now build the traded risk factors—the risk directions along which these shocks propagate. We use equilibrium returns and trading flows to capture all trades that intermediaries must absorb, such that our factors reveal the most important orthogonal, non-diversifiable risks intermediaries actually bear. Each factor's price sensitivity λ_k would then be estimated by regressing its return on its flow, instrumented with exogenous shifts in uninformed demand. Henceforth, we denote observed equilibrium trading flows by ΔQ and their instrumented counterparts—fitted values from exogenous shocks to uninformed demand—by $\Delta \hat{Q}$.

Conceptually, the factors we choose should (i) span the systematic risks generated by customer trading and (ii) remain free of cross-impacts. Achieving these objectives poses two difficulties. First, the complete set of systematic risks is unknown and potentially large (Cochrane, 2011). If we must work with a subset of factors, how can we ensure that this subset captures a substantial share of trading-induced risk, so that the chosen factors act as the primary channels through which demand shocks propagate across currencies? Second, under the standard equilibrium benchmark (2), uncorrelated returns are sufficient to eliminate cross-impacts. However, empirical evidence shows that correlated trading can also transmit shocks even when returns are orthogonal (Ben-David, Li, Rossi, and Song, 2022). How can we ensure that the constructed factors are truly orthogonal, with no residual cross-impacts?

Our "traded risk factors" address both challenges simultaneously. Using a new procedure applied jointly to return and flow data, we let customers' actual trading patterns reveal the risks they trade and care about, which plausibly channel demand propagation. We rank the factors by the share of trading-induced risk they explain, and select the top few knowing exactly how much risk this subset captures. Moreover, by construction, these factors exhibit both uncorrelated returns *and* uncorrelated trading flows, making them truly orthogonal.

The following proposition formalizes the idea and, together with Appendix A.2, provides an explicit algorithm for constructing the factors.

PROPOSITION 2. *There exists a set of K factors that spans all non-diversifiable risk induced by customer trading and satisfies*

1. *Uncorrelated factor returns:*

$$\text{cov}\big(r_k^{factor},\, r_j^{factor}\big) = 0,\ \textit{for any } k \neq j. \tag{17}$$

2. *Uncorrelated factor flows:*

$$\text{cov}\big(\Delta Q_k^{factor},\, \Delta Q_j^{factor}\big) = 0,\ \textit{for any } k \neq j. \tag{18}$$

The factors are ordered by descending $\text{var}(\Delta Q_k^{factor})\,\text{var}(r_k^{factor})$, *the amount of trading-induced risk each one explains.*

Proposition 2 is effectively a modified principal-component analysis (PCA) that operates *jointly* on asset returns and trading flows. By contrast, the standard return-only PCA is typically used to isolate factors that explain unconditional risk (Ross, 1976; Fama and French, 1993; Lustig, Roussanov, and Verdelhan, 2011). Concretely, the return-only PCA enforces (17), replaces (18) with orthogonal factor loadings, and sorts factors by descending unconditional return variance, $\text{var}(r_k^{\text{factor}})$. By incorporating both returns and flows, our procedure pinpoints the risk factors that customers trade most actively and quantifies each factor's contribution to trading-induced risks via the product $\text{var}(\Delta Q_k^{\text{factor}})\,\text{var}(r_k^{\text{factor}})$.

Our procedure also differs from a PCA applied solely to trading flows (Lo and Wang, 2000; Hasbrouck and Seppi, 2001; Balasubramaniam, Campbell, Ramadorai, and Ranish, 2023). The flow-only PCA enforces (18), replaces (17) with orthogonal portfolio weights, and sorts factors by descending unconditional quantity variance, $\text{var}(\Delta Q_k^{\text{factor}})$. When currency returns are independent and identically distributed, uncorrelated factor returns are equivalent to orthogonal weights because

$$0 = \text{cov}\big(r_k^{\text{factor}},\, r_j^{\text{factor}}\big) = \text{cov}\left(\sum_{n=1}^{N} w_{n,k} r_n,\ \sum_{n=1}^{N} w_{n,j} r_n\right) = \sum_{n=1}^{N} w_{n,k} w_{n,j}\text{var}(r_n). \tag{19}$$

In this case, the flow-only and our joint PCA coincide: all cross-sectional structure resides in the flow data. In general, combining flows with returns enables our method to recover the *risk factors* that customers actually trade, not merely the *portfolios* they trade. Section 4 empirically contrasts our approach with both return-only and flow-only PCA and documents the resulting differences.

3 Data

To identify traded risk factors, we need data on FX trading and returns. In this section, we outline the various data sources that we use.

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

We use daily aggregate FX order flow data from CLS, which includes the total value of buy and sell orders between Banks and their customers in 17 currencies from September 2012 to December 2023.[12] The currencies in our sample are: U.S. dollar (USD), Australian dollar (AUD), Canadian dollar (CAD), Swiss frank (CHF), Danish kroner (DKK), Euro (EUR), British pound (GBP), Hong Kong dollar (HKD), Israeli shekel (ISL), 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 involve Banks as one counterparty, where Banks include bank-affiliated dealers and hedge funds transacting through prime brokers. We interpret Banks' trading as representing the activities of the specialist intermediary in our model. Counterparties to Banks are grouped into three categories: Funds (e.g., mutual funds, pension funds, sovereign wealth funds), Non-bank Financials (e.g., insurance companies, clearing houses), and Corporates.

To measure the *total* FX risk borne by intermediaries, 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 pronounced negative correlation between flows into spot versus forward and swap, excluding either can underestimate the price sensitivity to risks (see Supplemental Appendix A). The CLS forward and swap data are organized by maturity buckets. We estimate FX spot exposure from these future-settled contracts by discounting the notional using forward rates.[13,14] Aggregating across

---

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

[12]Both the number of currencies covered and the sample start date are constrained by CLS data availability.

[13]Conceptually, FX swaps should not expose intermediaries to currency risk, as the spot and forward legs offset each other in notional amounts. However, a small amount of currency risk remains after discounting the forward leg. Our results are effectively unchanged if swaps are excluded.

[14]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 return at the weekly frequency. Weekly flows are calculated by summing daily flows from Thursday to the following Wednesday. Our final trading data is a panel spanning 2012-09-06 to 2023-12-31, consisting of weekly net inflow into 16 non-USD currencies, measured in USD, across spot, forward, and swap transactions.

3.2 Return Data

We obtain the forward and spot 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 result of borrowing USD at the US risk-free rate, converting to foreign currency at the spot exchange rate, earning the foreign risk-free rate, and converting back to USD at the future spot rate. For currency n from week t to $t+1$, we define $r_{t+1,n} = s_{t+1,n} - s_{t,n} + i_{t,n} - i_{t,USD} - x_{t,n} = s_{t+1,n} - f_{t,n}$, where s is the log spot rate, f is the log forward rate, i is the net risk-free rate, and x is the deviation from the covered interest-rate parity (CIP). Exchange rates are defined as USD per one unit of foreign currency, so a higher s corresponds to USD depreciation. Our currency return includes the CIP deviation, $x_{t,n} = f_{t,n} - s_{t,n} - i_{t,USD} + i_{t,n}$, to more accurately reflect the actual return that intermediaries have when absorbing customer flows, including inventory costs from balance sheet constraints.

3.3 Other Data

We collect sovereign bond auction data to instrument for FX demand shocks. Specifically, we source announcement information on auctions of bonds with maturities of one year or longer from government websites in the U.S., Australia, Canada, France, Germany, Italy, Japan, and the U.K.

To construct excess returns in five non-FX asset classes, we use the following data. For credit default swaps (CDS), we obtain five Markit indices from Bloomberg (North America investment grade and high yield, Europe main and crossover, and Emerging Market), with returns defined from the seller's perspective. For commodities, we use six Bloomberg commodity futures return indices (energy, grains, industrial metals, livestock, precious metal, and softs). For corporate bonds, we use four Bloomberg indices on U.S. corporate bonds by credit rating (Aa, A, Baa, high yield; excluding AAA to avoid collinearity with the risk-free

rate). For options, we calculate leverage-adjusted option portfolio returns on S&P 500 call and put prices from OptionMetrics, following Constantinides, Jackwerth, and Savov (2013). For US Treasury bonds, we use yields of the six maturity-sorted "Fama Bond Portfolios" from CRSP, excluding Treasury bills due to correlation with the risk-free rate. Finally, we use the 1-month U.S. Libor as a proxy for the risk-free rate.

The Bloomberg CDS data begin in 2007, OptionMetrics data end in December 2022, and all other asset classes data span January 2000 to December 2023.

4   Traded Risk Factors in FX

In this section, we identify the most traded FX factors from data. We first find that three risk factors account for most of the non-diversifiable risks induced by FX trading. We then interpret these factors as the Dollar, the Carry, and the Euro-Yen Residual. Finally, we show that these factors cannot be obtained by the standard PCA on returns or flows alone.

4.1   Baseline Traded FX Factors

Our objective is to identify risk factors that capture the effect of FX trading on currency prices in the cross-section. To this end, we focus on factors that maximally explain trading-induced risks. Using the procedure detailed in Section 2.4, we derive the traded FX factors from weekly net flows and log returns of 16 non-USD currencies.[15] The three factors that explain the most amount of trading-induced risk are reported in Table 1. Each column of Table 1 represents a factor, and the component values 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.[16] Because the identified risk factors are traded, they place greater weight on widely traded currencies. Notably, six developed economy currencies — AUD, CAD, CHF, EUR, GBP, and JPY — have consistently high weights across the top three factors; they are highlighted in red along with USD. Of the total trading-induced non-diversifiable risks, the top three traded FX factors individually account for 65%, 16%, and 9%, respectively. Jointly, these three factors explain approximately 90% of the risks intermediaries bear when accommodating trading flows. The traded FX factors are stable over time. Table SA3 in the Supplemental Appendix shows that the returns and flows of factors recovered from the

---

[15]We use aggregate flows across all customer types to identify total trading-induced risks from the inter-mediaries' perspective. Trades from different customers may carry different informational content but pose the same balance-sheet or inventory risk.

[16]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 3 Traded FX 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	0.02
EUR	-0.5	-0.43	1.16
GBP	-0.11	0.18	0.09
HKD	0	-0.01	0.02
ILS	0	0	0
JPY	-0.07	-0.49	-1
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	0.02
ZAR	-0.01	0.01	-0.01
USD	1	0.03	0.74
Var explained	65%	16%	9%

Notes: This table presents the portfolio weights of the top 3 traded FX factors, constructed following the procedure in Section 2.4. The return and 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.

full sample are highly correlated with those recovered from the pre-2020 or the post-2020 subsamples. Indeed, the correlations are nearing 1 for the first factor and exceeding 0.8 for the other two.

4.2 Interpretation of Traded FX Factors

To better understand the risks captured, we conjecture and verify that the top three traded FX factors represent the Dollar, the Carry, and the Euro-Yen Residual, respectively. Factor 1 in Table 1 assigns negative weights to all non-USD currencies, resembling the proverbial Dollar portfolio that shorts all non-USD currencies to bet on the USD exchange rate. Appendix A.3 shows that our factor-construction algorithm is invariant to the choice of numeraire. The Dollar's emergence as the first factor therefore reflects its economic status as the world's reserve currency—not a mechanical consequence of expressing returns in USD. We therefore propose a traded Dollar factor that goes long in USD and shorts the six most

Table 2: **Correlation between Return and Flow for Baseline PC Factors versus for Proposed Economic Factors**

	Factor 1	Factor 2	Factor 3
Return	0.98	0.95	0.92
Flow	1.00	0.99	0.95
Var explained by Economic Factors	63%	15%	8%

Notes: This table shows the correlation between return and flow for the baseline traded FX factors in Table 1 ("PC Factors") and for the traded FX factors constructed from the proposed factor weights of the Dollar, the Carry, and the Euro-Yen Residual ("Economic Factors"). It also shows the fraction of trading-induced risks explained by the Economic Factors.

traded currencies (AUD, CAD, CHF, EUR, GBP, and JPY) in equal weights. Factor 2 has positive weights on high-interest-rate currencies (e.g., AUD, CAD, GBP) and negative weights on low-interest-rate currencies (e.g., JPY, CHF, EUR), consistent with the proverbial Carry portfolio that exploits violations of uncovered interest-rate parity (UIP). We propose a traded Carry factor that goes long in 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 a traded Euro-Yen Residual factor that goes long in EUR and shorts JPY in equal weights. The rationale is that, because EUR and JPY are traded in the same direction in both Dollar and Carry factors, these factors do not capture the bilateral trading flows between the Euro area and Japan, two of the world's largest economies.

These proposed factors are economically meaningful but may be correlated. To address this, we apply the procedure described in Section 2.4 to orthogonalize them. In particular, this process transforms the proposed EUR-JPY pair (long EUR, short JPY) into the Euro-Yen Residual factor, which is uncorrelated with the Dollar and Carry factors. In other words, the Euro-Yen Residual factor captures the portion of non-diversifiable risk that intermediaries bear when absorbing EUR-JPY pair trading, after hedging out exposures to 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 our interpretation of the traded FX factors. Table 2 shows the correlation between returns and flows of the baseline factors ("PC Factors") from Table 1 and returns and flows of the factors constructed from the proposed Dollar, Carry, and Euro-Yen

Residual weights ("Economic Factors"). The correlations are nearly 1 for both returns and flows across all three factors. Together, the three Economic Factors explain about 86% of trading-induced non-diversifiable risks, closely matching the risks accounted for by the PC Factors. Given this striking similarity and to avoid potential in-sample overfitting concerns with PC Factors, we focus on analyzing the more interpretable Economic Factors for the remainder of the paper.

Panel (a) of Figure 1 plots the cumulative trading flows from customers to the three traded FX factors.[17] During our sample period, customers purchased approximately $1 trillion of the Dollar factor from intermediaries, primarily after the 2020 COVID crisis. This provision of USD by intermediaries likely reflects USD deposits or wholesale funding made available by (dealer-affiliated) banks (Du and Huber, 2024), as some intermediaries, especially dealers, may not be able to maintain a sustained inventory imbalance. For the Carry factors, customers initially refrained from large directional bets but began selling off the Carry factor post-2022. As a result, intermediaries including dealers and hedge funds accumulated $0.8 trillion in Carry trade exposure between 2012 and 2023. Finally, customers sold the Euro-Yen Residual factor up until the 2020 COVID crisis, after which they started repurchasing some, but not all, positions. This left the intermediaries with a net positive position in the Euro-Yen Residual factor throughout the sample period. As JPY acts as a "funding currency" (negative weight) in both the Carry and Euro-Yen Residual factors, our analysis highlights that the unwinding of intermediaries' short JPY positions cannot solely be attributed to the Carry trade.

Panel (b) of Figure 1 plots the cumulative returns of the three factors over our sample period. We observe that all three factors enjoy positive returns, including the Euro-Yen Residual factor. We formally investigate the unconditional risk premium of these factors in Section 5.1.

4.3 Standard PCA on Returns or Flows Fails to Identify Traded Risk Factors

We demonstrate that a standard PCA applied solely to returns or flows fails to identify the traded FX factors. The results underscore the empirical value of using our approach to jointly analyze returns and flows.

The first three columns of Table 3 show the portfolio weights for the first three principal components of a standard PCA applied to returns.[18] The first factor resembles a Dollar factor, with negative loadings on all currencies. The second factor assigns large positive

---

[17]Figure SA2 in the Supplemental Appendix provides a breakdown of factor flows by customer type.

[18]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).

Figure 1: **Cumulative Flow and Return of Top 3 Traded FX Factors**


(a) **Cumulative Flow**

(b) **Cumulative Return**

Notes: This figure displays the cumulative flows and returns of the top three traded FX factors between September 2012 and December 2023. Flows are measured from the perspective of customer purchases (intermediary sales). For instance, the figure indicates that customers bought approximately $1 trillion of the Dollar factor from intermediaries during this period.

weights to some high-interest-rate currencies such as ZAR and MXN, and large negative weights to some low-interest-rate currencies like CHF, JPY, and EUR. However, it also assigns very small positive weights to other high-interest-rate currencies like AUD and NZD and even a negative weight to GBP and NOK.[19] 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

---

[19]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.

Table 3: **Top 3 PCs from FX Returns or Flows**

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	-0.06
CHF	-0.05	-0.21	-0.51	-0.01	-0.02	-0.06
DKK	-0.06	-0.15	-0.12	0	0	0.01
EUR	-0.06	-0.15	-0.13	-1	-0.03	0.03
GBP	-0.07	-0.08	0.47	-0.02	-0.01	0.26
HKD	0	0	0	0	-0.02	0
ILS	-0.04	-0.03	0.24	0	-0.01	0
JPY	-0.03	-0.17	-1	-0.04	-0.06	-0.95
KRW	-0.06	0.02	-0.15	0	0.01	0
MXN	-0.08	0.22	0.71	-0.01	0.01	0
NOK	-0.1	-0.05	0.72	0	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	0
SGD	-0.04	-0.03	-0.12	-0.01	-0.01	0.01
ZAR	-0.11	0.29	-1.35	-0.01	0	0.01
USD	1	0.37	0.29	1.17	-0.92	0.62

Notes: The first three columns display the portfolio weights for the first three principal components from a return PCA, while the second three columns show those from a flow PCA. The analysis uses weekly data for 16 non-USD currencies spanning September 2012 to December 2023. The USD portfolio weight is calculated as the negative sum of the weights of all other currencies.

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 flows. The resulting portfolios from this approach 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, entirely overlooking the strong factor structure in returns.

Table 4: **Unconditional Risk Premium**

	Dollar	Carry	Euro-Yen Residual
Panel A: Sep 2012 to Dec 2023			
Mean return (annualized %)	2.38	2.15	5.26
Sharpe ratio (annualized)	0.35	0.26	0.56
Fama-MacBeth premium (annualized %)	2.42	3.34	3.58
t-stats	(1.15)	(1.22)	(1.12)
Panel B: Jan 2000 to Dec 2023			
Mean return (annualized %)	-0.16	2.09	1.99
Sharpe ratio (annualized)	-0.02	0.23	0.20
Fama-MacBeth premium (annualized %)	-0.07	3.02	1.00
t-stats	(-0.04)	(1.41)	(0.40)

Notes: This table presents the annualized mean return and Sharpe ratio of the three traded FX factors. Additionally, it reports the Fama-MacBeth factor premium along with t-statistics calculated using Shanken-corrected standard errors. Panel A is based on weekly returns from September 2012 to December 2023, while Panel B uses weekly returns from January 2000 to December 2023.

5 Pricing Properties of Traded FX Factors

In this section, we study the traded FX factors' unconditional risk premium and their price sensitivity to trading-induced risks.

5.1 Unconditional Risk Premium

Panel A of Table 4 reports the annualized mean returns and Sharpe ratios of the three traded FX factors based on weekly returns from September 2012 to December 2023. Notably, the newly proposed Euro-Yen Residual factor achieves an annualized return exceeding 5% and a Sharpe ratio of 0.56, both meaningfully higher than those of the other two factors. To evaluate the cross-sectional pricing power of these factors, we estimate the Fama-MacBeth factor premia.[20] The Fama-MacBeth premia of the three factors are similar to their mean returns estimated from the time series, though we caution that the estimated Fama-MacBeth premia are not statistically significant, which may partly reflect that the portfolios are static and not conditionally rebalanced as in Lustig, Roussanov, and Verdelhan (2011).

Our sample period begins in September 2012 due to the availability of CLS data. To

---

[20]We follow the Fama-MacBeth two-step procedure: first, time-series regressions of each currency's return on factor returns estimate betas; second, cross-sectional regressions of average currency returns on these betas (excluding the constant) recover the factor premium. Standard errors are corrected following Shanken (1992).

further explore unconditional risk premia, we extend the sample to start in 2000 (introduction of the Euro) and report the results in Panel B. In this longer sample, the Euro-Yen Residual factor exhibits a time-series mean return and Sharpe ratio comparable to the Carry factor. In the cross-section, the Carry factor demonstrates considerably stronger pricing power than the other two factors.

5.2 Price Sensitivity to Trading-Induced Risks

We aim to estimate λ_k, the price sensitivity to trading-induced risks of traded FX factor k in equation (11). Because the traded FX factors are constructed to have uncorrelated returns and uncorrelated flows, the theoretical foundation ensures that λ_k can be estimated factor-by-factor without concern for cross-factor substitution. However, for each factor, we must instrument for the unobserved demand shocks that are unrelated to changes in fundamentals. Specifically, we regress each factor's risk-adjusted returns[21] on its instrumented weekly flows, $\hat{q}_{k,t}$:

$$r_{k,t}^{\text{factor}}/\text{var}(r_{k,t}^{\text{factor}}) = \lambda_k \Delta \hat{Q}_{k,t}^{\text{factor}} + \epsilon_{k,t}, \text{ where} \tag{20}$$

$$\Delta Q_{k,t}^{\text{factor}} = \theta_k z_{k,t} + e_{k,t}, \tag{21}$$

$$\text{cov}(z_{k,t}, \epsilon_{k,t}) = 0. \tag{22}$$

The instruments (z_k) for the observed factor flows $(\Delta Q_{k,t}^{\text{factor}})$ must be both relevant (equation (21)) and valid (equation (22)). We propose sovereign bond auction *announcements* as instruments.[22] Government entities, such as the U.S. Treasury, periodically auction off long-term debt obligations, e.g., U.S. Treasury notes and bonds. 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.[23] When auctions are announced about a week in advance, these announcements can prompt foreign investors to exchange domestic currencies for local currencies, making these instruments relevant.

We also argue that the instruments are valid. First, auction announcements are plausibly

---

[21]Each factor's weekly observed return $r_{k,t}^{\text{factor}}$ is normalized by its annualized return variance, $\text{var}(r_{k,t}^{\text{factor}})$, so the regression coefficient estimates the price sensitivity to risk λ_k, as defined in equation (11). Technically, theory implies that the return variance should be measured in the absence of demand shocks. In practice, we use realized variance because the demand shocks we study are small in magnitude and likely contribute little to the overall return variance.

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

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

exogenous to FX trading because auctions follow strong fiscal cyclicality and are largely predetermined. For example, the U.S. Treasury Borrowing Advisory Committee (TBAC) issues two-quarter-ahead recommendations on debt issuance for upcoming auctions. The subsequent announcements and the eventual issuance of long-term debt (maturities of longer than a year) rarely deviate from these recommendations (Rigon, 2024).[24] Although these auction-induced demands are predictable, they can still move prices when they materialize. Theoretically, sophisticated investors do not find it optimal to front-run the entire expected demand (Vayanos, 2021); empirically, demand shocks generated by anticipate stock dividend payments nonetheless move stock prices (Hartzmark and Solomon, 2024).

Second, auction announcements plausibly satisfy the exclusion restriction that their effect on exchange rates arises solely through FX trading. Because auctions are heavily forward-guided, announcements likely contain limited new information that would affect fundamentals ($\epsilon_{k,t}$). Another concern is that auction announcements might induce excess bond trading, affecting bond prices and spilling over to FX. However, empirically, Wachtel and Young (1990) find that while Treasury auction *results* move bond yield, the week-ahead *announcements* have no detectable effect. Thus, any impact on FX is likely driven solely by announcement-induced FX demand shocks.

Finally, the behavior of asset prices provides a test for the validity of our instrument. As discussed in Section 2, valid demand shocks move currency prices initially (at time 0), but these price effects should eventually revert (at time 1). In contrast, if the considered demand shocks contained new information, the resulting price response would be permanent and not revert. Empirically, Figure SA3 in the Supplemental Appendix shows that the contemporaneous price responses of all three factors fully revert within a month.

As the traded FX factors place weights on multiple currencies, we consider sovereign auction announcements from a panel of countries. Specifically, U.S. Treasury auction announcements instrument demand shocks to the Dollar factor; Australian, Canadian, British, and Japanese government bond auction announcements instrument shocks to the Carry factor, and Euro-Area government bond auctions (aggregating German, French, and Italian auctions) instrument for the Euro-Yen factor. For each factor, we aggregate the offered amount across all announcements in a week, consistent with FX trading flows.[25] Finally, we remove any linear trend in auction sizes over time.

---

[24]Similarly, Germany's Finance Agency releases an annual auction calendar each December, specifying target amounts for each auction.

[25]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 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 5: **Estimated Price Sensitivity to Trading-Induced Risks**

| | Dollar | | Carry | | Euro-Yen Residual | |
| | OLS | IV | OLS | IV | OLS | IV |
	(1)	(2)	(3)	(4)	(5)	(6)
Factor flow	0.072***	0.107***	0.132***	0.138**	0.139***	0.335*
	(0.009)	(0.037)	(0.018)	(0.064)	(0.021)	(0.195)
Response per \$B (bps)	3.4	5.0	8.9	9.3	12.2	29.3
1st stage F-stat		24.8		6.5		3.8
Anderson-Rubin CI				(0.01, 2.39)		(0.09, 1.91)
Observations	590	386	590	228	590	560

Notes: This table presents the λ_k estimation results for the Dollar, Carry, and Euro-Yen Residual factors, based on regression (20). The response of factor prices to demand shocks, measured per billion dollars, is calculated as the product of λ_k and the annualized return variance. The IV regressions report the first-stage heteroscedasticity and autocorrelation consistent (HAC) effective F-statistics and the Anderson-Rubin confidence intervals at the 90% confidence level. 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. *$p <.1$; **$p <.05$; ***$p <.01$.

Table 5 presents the λ_k estimation results for the Dollar, Carry, and Euro-Yen Residual factors. For all three factors, the estimated price sensitivity to trading-induced risks is positive and statistically significant. Recall that the regression (20) normalizes each factor's return by its variance. As a result, the estimated λ_k captures the price response to one unit of *risk* induced by \$1 billion of factor flow and is directly comparable across factors. Both OLS and IV estimates show that the price sensitivity to risks is the smallest for the Dollar, higher for the Carry, and highest for the Euro-Yen Residual. This indicates that intermediaries bear marginal risks most effectively in the Dollar factor, with their risk-bearing capacity progressively lower for the Carry and the Euro-Yen Residual. Viewed through the model in Section 2.2, the cross-factor variation in price sensitivity to risks may reflect differences in available arbitrage capital across risk factors, with lesser-known factors like Euro-Yen Residual attracting less arbitrage capital. Indeed, the annualized volatility of trading flows absorbed by intermediaries is highest for the Dollar factor (\$85 billion), followed by the Carry factor (\$ 34 billion), and lowest for the Euro-Yen Residual factor (\$22 billion). The OLS estimates are slightly smaller than the IV estimates, reflecting the instrument's role in mitigating bias from the correlation between information-driven price changes $\epsilon_{k,t}$ and contemporaneous customer flows $\Delta Q_{k,t}^{\text{factor}}$. This correlation is negative, likely because customers trade against fundamentals: they buy when news causes a currency to depreciate and sell when it appreciates. Such behavior is consistent with the profitability of momentum

strategies in FX (Menkhoff, Sarno, Schmeling, and Schrimpf, 2012).[26]

To compare the magnitude of our estimated price sensitivity to risks with the literature, we multiply each factor's λ_k by its return variance to calculate the factor-level price response per billion of demand shocks, as shown in the second row of Table 5. A \$1 billion demand shock increases the prices of the Dollar, Carry, and Euro-Yen factors by 5, 9, and 29 basis points, respectively.[27] These price responses are large compared to U.S. equities, where a \$1 billion demand shock to the entire U.S. stock market raises the aggregate price by about 1.7 bps (Gabaix and Koijen, 2021).[28] We think that the supply of FX arbitrage capital is likely limited due to the specialized nature of the FX market, where only sophisticated participants like bank dealers and hedge funds absorb demand shocks.[29] This may seem counterintuitive given the large turnover in FX, but up to 75% of trades occur between intermediaries (BIS, 2022), suggesting that the arbitrage capital available to absorb shocks is much smaller than the total turnover.[30]

Because our estimates are obtained through IV analysis, which reflect the local average treatment effect, we caution against extrapolating our estimates to much larger demand shocks, such as quantitative easing. Nevertheless, our results can be informative about the effect of larger interventions. Specifically, as our estimates are based on relatively small, high-frequency shocks, intermediaries likely absorb these shocks with a risk-bearing capacity that is fixed in the short-run. If there are larger shocks, however, intermediaries may expand their capacity (e.g., bank-affiliated dealers could tap into the banks' broader balance sheets). Such global re-optimization likely leads to a smaller average price response per unit of risk, as the price impact function may be concave (Hasbrouck, 1991; Chaudhry and Li, 2025). Our estimates thus form an upper bound for the price impact of larger shocks.

Additionally, because the price impacts of these factor-level demand shocks revert within a month, Appendix B derives an expression for the share of return variance contributed by

---

[26]In a rational market, prices would adjust to fundamental news without trading (Milgrom and Stokey, 1982). However, when customers buy in response to negative fundamental news, prices under-react, leading to subsequent price drift and generating momentum.

[27]In a dynamic setting, the persistence of demand shocks can influence price response, as intermediaries anticipate future demand (e.g., Campbell and Kyle, 1993; Wang, 1993; Jansen, Li, and Schmid, 2024). Our estimates reflect the average level of persistence over the sample period.

[28]Gabaix and Koijen (2021) find that a 1% greater demand shock to the entire US stock market increases price by 5%. Given an average market capitalization of \$31.7 trillion between 2012 and 2022, a \$1 billion demand shock raises the price of the market factor by 1.7 bps over our sample period.

[29]The limited FX arbitrage capital may also reflect slow-moving capital and the fact that our price sensitivity to risks is estimated based on a weekly horizon, shorter than the monthly or quarterly horizons typically considered in the literature. Asset markets tend to be more inelastic over shorter horizons as long-term investors are slower to react to price changes and provide arbitrage capital (Duffie, 2010).

[30]Of the FX trades accounted for in the BIS Triennial Central Bank Survey, 46% are between reporting dealers, 22% with non-reporting dealers, and 7% with hedge funds, all of which are intermediaries in our model and captured in Banks in the data.

trading volatilities over different horizons. Applying this formula across the three factors, we find that flows account for about 10–35% of variance at the 1-week horizon, 5–15% at one month, and fade quickly thereafter. These results confirm that demand shocks are a meaningful yet short-lived source of return volatility.

Finally, the precision of IV estimation depends on the strength of the instrument. The heteroscedasticity and autocorrelation consistent (HAC) effective F-statistics for the Dollar, the Carry, and the Euro-Yen factors are 24.8, 6.5, and 3.8, respectively. The effective F-statistics for the Carry and the Euro-Yen are below the rule-of-the-thumb threshold of 10. To assess the implications of potentially weak instruments on IV inference, we compute the Anderson-Rubin confidence interval, which has the correct coverage regardless of the strength of the instrument (Andrews, Stock, and Sun, 2019). For both the Carry and the Euro-Yen Residual, the Anderson-Rubin confidence interval is bounded away from zero, but is very wide in the positive direction. In other words, we are reasonably confident that the price sensitivity to risks is not zero but much less certain that the true value is not larger. A larger estimate would mean an even greater price sensitivity to risk.

5.3   Time-Varying λ and the Role of Risk

Our representative intermediary framework posits that price responses to trading stem from intermediaries' sensitivity to risk. In the previous subsection, we discussed patterns in the estimated λ consistent with this view—for instance, specialization may limit arbitrage capital and risk-bearing capacity, resulting in larger price responses. In this subsection, we seek more direct evidence that risk drives observed price responses to trading. Specifically, we examine whether λ depends on time-varying wealth or constraints that alter intermediaries' risk-return trade-off.

We consider two proxies. First, we use intermediary equity returns to capture intermediaries' wealth.[31] Second, we use deviations from covered interest-rate parity (CIP) to capture intermediaries' constraints, as such deviations indicate intermediaries' inability to exploit known profitable trades.[32]

Table 6 presents potential determinants of the Dollar factor's weekly return. We emphasize the state-dependency of the Dollar factor's λ, as the Dollar is the most traded FX factor and its flow instrument exhibits the highest statistical power. Column (1) suggests

---

[31]Following He, Kelly, and Manela (2017), we construct the value-weighted weekly return of primary dealers' bank holding companies. This series is highly correlated (0.95) with the KBW NASDAQ bank index over our sample period and is equivalent to the intermediary capital ratio shock (0.98 correlation) in He, Kelly, and Manela (2017).

[32]We calculate the weekly average cross-currency basis using the AUD-JPY currency pair and 3-month IBOR.

Table 6: **Time-Varying λ for the Dollar Factor**

	Weekly Return of Dollar Factor					
	(1)	(2)	(3)	(4)	(5)	(6)
Intermed. ret	-0.490***	-0.109				
	(0.119)	(0.204)				
Flow × Intermed. ret		-0.091***				
		(0.033)				
S&P ret			-0.148	-0.077		
			(0.096)	(0.314)		
Flow × S&P ret				0.006		
				(0.074)		
CIP deviation					0.081	0.182
					(0.060)	(0.177)
Flow × CIP deviation						0.063
						(0.129)
Factor flow		0.096***		0.106***		0.160*
		(0.037)		(0.040)		(0.093)
Observations	559	385	559	385	559	385

Notes: This table reports the IV-estimated time-varying λ for the Dollar factor. "Interm. ret" is the value-weighted weekly equity return of primary dealers' bank holding company. "S&P ret" is the weekly return of the S&P 500 index. "CIP deviation" is measured by the weekly average AUD-JPY 3-month IBOR cross-currency basis. All three variables are demeaned and standardized. All factor flows are instrumented with U.S. Treasury auction announcements. Newey-West standard errors are reported in parentheses, where the bandwidth is chosen by the Newey and West (1994) selection procedure. $*p <.1$; $**p <.05$; $***p <.01$.

that the Dollar factor reflects variations in intermediary equity returns. However, Column (2) clarifies that intermediary equity returns do not directly affect the Dollar's return. Instead, they influence λ, consistent with a risk-based view of price response: as intermediaries' wealth increases, their effective risk aversion decreases, reducing the price response to absorbing demand shocks (instrumented using U.S. Treasury auction announcements). This state-dependent response is driven specifically by intermediaries' wealth, as Columns (3) and (4) show that broader stock market returns have no comparable effect on λ. Conceptually, intermediaries' constraints may also affect price response: when constraints prevent intermediaries from fully exploiting profitable investment opportunities, they become more selective, leading to higher effective risk aversion and lower risk-bearing capacity. Empirically, the effects of such constraints, proxied by CIP deviations, are directionally consistent with the risk-return trade-off but, as shown in Column (6), not statistically significant.

Figure 2: **Decomposition of Currency Returns Explained by Traded FX Factors**

Notes: This figure plots the R^2 of regressing currency-level returns against the returns of the Dollar, the Carry, and the 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.

6   Demand Propagation Across Currencies and Asset Classes

In this section, we use the traded FX factors' estimated price sensitivity to risks to study the propagation of demand shocks among currencies and asset classes. We quantify demand propagation with cross-multipliers: the effect of a shock to demand for one asset on the price of another, holding all other demand constant.

6.1   Demand Propagation Across Currencies

For a traded FX factor to affect currency-level cross-multipliers, the currencies must load on the factor. Figure 2 demonstrates the relevance of the traded FX factors in explaining individual currency returns. Regressing currency-level returns on the returns of the Dollar, the Carry, and the Euro-Yen Residual factors in the time series, we plot the marginal R^2 attributed to each factor, which are additive because the factor returns are orthogonal by construction. The positive and negative signs in the plot indicate the direction of each currency's beta loading on each factor. Together, the three factors explain between 69% and 94% of individual currency returns.

The decomposition in Figure 2 provides a framework to analyze the risk implied in demand shocks. For instance, when a customer buys $1 of AUD from intermediaries, Figure 2 shows that intermediaries attribute 60% of the total risk to the Dollar factor, 10% each to the Carry and Euro-Yen Residual factors, and 20% to idiosyncratic risk unexplained by the

three factors. The direction of factor loadings further reveals that intermediaries perceive the customer's \$1 purchase (and their \$1 sale) of AUD as the customer selling the Dollar and Euro-Yen Residual factors while buying the Carry factor.

Combining the information in Figure 2 with the IV estimated price sensitivity to risks λ_k, we compute the cross-currency multipliers according to Proposition 1 and report the results in Table 7. For clarity, we have arranged the six major currencies (AUD, CAD, GBP, CHF, EUR, JPY) in the upper left quadrant, followed by the other ten currencies in the sample. Each entry shows the price response in one row (column) currency, in basis points, to a \$1 billion demand shock to the corresponding column (row) currency. For instance, the entry of 7.9 in the first row and second column indicates that a \$1 billion demand shock to the CAD (AUD) raises the price of AUD (CAD) by 7.9 bps (in percentage terms), holding the demand for all other currencies equal. Because the model-implied cross-multiplier is symmetric, we report only the upper half. The diagonal entries represent each currency's own multiplier. As discussed in relation to the literature, the diagonal multipliers in Table 7 reflect price responses due to the three most traded FX factors; they exclude price impacts arising from changes in a currency's idiosyncratic risk.

Table 7 reveals several interesting patterns of cross-currency multipliers. First, all entries are positive. This is because all currencies load on the Dollar factor in the same direction, which is the most traded risk factor in the cross-section. Second, the cross-multiplier between currencies on the long leg of the Carry trade (e.g., AUD, CAD, GBP) and those on the short leg (e.g., CHF, EUR, JPY) is generally smaller. The modest cross-multipliers reflect opposite beta loadings with respect to the Carry factor, which makes currencies in one group effective hedges for the Carry risk exposure of the other group. In short, these two groups are "complements" in their exposure to the Carry risk factor. Third, we note that although EUR and JPY are both low-interest-rate currencies, we estimate a rather small cross-multiplier because the two currencies are on the opposite side of the Euro-Yen Residual factor. This result suggests that EUR and JPY are not entirely substitutable.

Moreover, although we analyze traded FX factors constructed based on the six major currencies and USD, we recover meaningful cross-multiplier in other currencies due to these currencies' loadings on the three traded FX factors. Finally, as a sanity check of our methodology, we examine the cross-multiplier for HKD, a currency pegged to USD within a narrow band of 1%. While we do not use this pegged information in our estimation, the estimated cross-multipliers in the entire column and row associated with HKD are close to zero. This minimal impact reflects the nature of a pegged currency: its own demand shocks have negligible risk implications for other currencies, and its exchange rate relative to USD is largely unaffected by demand shocks to other currencies.

Table 7: **Demand Propagation Across Currencies Through the Three Most Traded FX Factors**

	AUD	CAD	GBP	CHF	EUR	JPY	DKK	HKD	ILS	KRW	MXN	NOK	NZD	SEK	SGD	ZAR
AUD	12.0	7.9	9.0	2.1	2.8	5.9	2.8	0.2	4.7	6.3	7.8	10.4	10.4	5.9	4.3	11.0
CAD		5.3	5.9	0.7	1.6	2.6	1.6	0.1	3.0	4.0	5.3	6.8	6.7	3.7	2.6	7.2
GBP			7.4	3.1	4.0	3.2	3.9	0.1	3.9	5.0	6.2	8.9	8.0	6.1	3.5	8.8
CHF				8.6	7.3	4.1	7.3	0.0	2.4	2.4	1.1	5.1	2.7	6.5	2.4	3.2
EUR					7.4	0.2	7.4	0.1	2.5	2.4	2.5	6.1	3.1	7.1	2.3	4.2
JPY						16.2	0.2	0.0	2.3	4.0	0.9	3.5	5.7	1.1	3.1	4.0
DKK							7.4	0.1	2.5	2.4	2.5	6.0	3.1	7.1	2.3	4.2
HKD								0.0	0.1	0.1	0.1	0.2	0.1	0.1	0.1	0.2
ILS									2.1	2.7	3.1	4.8	4.2	3.5	2.0	4.6
KRW										3.6	4.0	5.9	5.6	3.9	2.5	5.9
MXN											5.7	7.3	6.6	4.6	2.6	7.5
NOK												11.1	9.4	8.2	4.3	10.5
NZD													9.1	5.7	3.9	9.6
SEK														7.7	3.1	6.8
SGD															1.9	4.1
ZAR																10.5

Notes: This table uses Proposition 1, the estimated factor-level price sensitivity to risks λ_k from Table 5, and the beta loadings of currencies to factors (signs illustrated in Figure 2) to compute currency-level cross-multiplier. Each entry represents the percentage price change in bps of a row (column) currency, as induced by a \$1 billion demand shock to a column (row) currency, holding the demand in all other currencies equal. As noted after Proposition 1, the model-implied cross-multiplier is symmetric, so we report only the upper half.

32

Figure 3: **Decomposition of Asset-Class Returns Explained by Traded FX Factors**

Notes: This figure plots the R^2 of regressing individual asset's monthly excess returns against the returns of the Dollar, the Carry, and the 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. The estimation period is from 2000-02 to 2023-12. The returns from CDS are available starting in 2007-04. The returns from Opt end in 2022-12.

6.2 Demand Propagation Across Asset Classes

If other asset classes load on the traded FX factors, demand shocks can propagate through shared FX exposures. We analyze five non-FX asset classes: credit default swap (CDS), commodities (Comm), corporate bonds (CorpBond), options (Opt), and US Treasury bonds (UST).[33] Similar to Figure 2, we regress the monthly excess returns of each asset class from 2000-02 to 2023-12 on the Dollar, Carry, and Euro-Yen returns, and present the R^2 decomposition in Figure 3.[34,35]

The three traded FX factors jointly explain between 15% (commodities) and 37% (corporate bonds) of the returns in the five non-FX asset classes we examine. The high explanatory power of traded FX factors is not an artifact of crisis-period comovements. As shown in Figure SA4 of the Supplemental Appendix, results based on returns excluding the 2007–09 Financial Crisis and the COVID-19 period are largely similar. Interestingly, while the Dollar

---

[33]We exclude equities because Haddad and Muir (2021) show that intermediation in equities differs considerably from FX, suggesting different marginal investors. While traded FX factors may partially explain equities returns, their price sensitivities are likely different in equities.

[34]We construct the return of each asset class as the equal-weighted average return of all available portfolios; see also Section 3.3.

[35]By construction, the correlation among weekly factor returns is zero. The correlation among monthly factor returns is close to zero. We report the incremental R^2 by adding the factors sequentially in the order of the Dollar, the Carry, and the Euro-Yen.

Table 8: **Demand Propagation Across Asset Classes Through Shared FX Risks**

	CDS	Comm	CorpBond	Opt	UST
CDS	2.4	3.5	3.2	4.7	-0.5
Comm		8.9	6.0	7.7	0.7
CorpBond			4.8	6.5	-0.2
Opt				9.3	-0.6
UST					0.7

Notes: This table uses Proposition 1, the estimated factor-level price sensitivity to risks λ_k from Table 5, and the beta loadings of assets to factors (signs illustrated in Figure 3) to compute asset-level cross-multiplier. Each entry represents the percentage price change in bps of a row (column) asset, as induced by a \$1 billion demand shock to a column (row) asset, holding the demand for all other assets equal. As noted after Proposition 1, the model-implied cross-multiplier is symmetric, so we report only the upper half.

factor is statistically significant across all five asset classes, it is least important in explaining the return of U.S. Treasury bonds (Treasurys).[36] Moreover, while all other asset classes load positively on the Carry factor, Treasurys load negatively. This contrast suggests that large shocks to the Carry factor could drive divergent price movements between Treasurys and other assets. Finally, while the Euro-Yen Residual factor is less prominent in non-FX asset classes, it explains a non-negligible fraction of returns in corporate bonds.

Similar to Table 7, we report cross-multipliers between asset classes in Table 8.[37] Examining the last column of Table 8, we recover two salient features of Treasurys while using only assets' factor loadings and factors' price sensitivity to risks. First, the price response to a demand shock is smallest for Treasurys, corroborating the observation that the Treasury market is deep and liquid. Second, Treasurys uniquely exhibit negative cross-multipliers with most other asset classes. A \$1 billion demand shock to Treasurys raises their price but depresses the price of other assets, reflecting Treasurys' "safe haven" property. Our estimation captures this behavior because only Treasurys load negatively on the commonly priced Carry factor, making them an effective hedge against other asset classes during shifts between "risk-on" and "risk-off" regimes.

We raise two cautions in interpreting our estimated cross-asset multipliers. First, our estimates capture only demand propagation across asset classes through exposure to the three traded FX factors. They may not represent the total price response to a \$1 demand

---

[36]One possible reason for this attenuated connection is that foreign investors hedge a substantial amount of the USD FX risks associated with their securities holdings, especially bonds (Du and Huber, 2024).

[37]The cross-multiplier between the traded FX factors and these five non-FX asset classes are reported in Table SA4 of the Supplemental Appendix.

shock to an asset, as these assets may also be exposed to other shared risks that we do not capture. Second, by using λ_k from the traded FX factors to inform multipliers in other asset markets, our analysis implicitly assumes that the marginal intermediaries are the same across different markets. Departures from this assumption may alter the magnitude but not the mechanism of demand propagation.

7   Conclusion

In conclusion, this paper studies the propagation of demand shocks through traded risk factors. If asset prices respond to risks and marginal investors can diversify risks across assets, then demand shocks propagate by affecting non-diversifiable risks, as captured by traded risk factors. We identify the most traded risk factors by extending the concept of priced non-diversifiable risks (Ross, 1976) to a representative intermediary framework (He and Krishnamurthy, 2017) and developing a method that integrates trading and returns data.

Applying the method to FX, we uncover the Dollar, the Carry, and the Euro-Yen Residual factors. These three factors explain 90% of the non-diversifiable risk intermediaries absorb in FX trading, and IV estimates show factor prices rise by 5–30 basis points per \$1 billion of factor demand. Hence, a demand shock to one currency propagates by changing the demand for these traded risk factors, affecting first the factors' prices and then the price of other currencies with shared exposures. The same logic leads to demand propagation across non-currency assets through shared currency risk. In short, we link trading quantities and asset prices (Froot and Ramadorai, 2008; Koijen and Yogo, 2019) through risks, underscoring the role of common risk in cross-asset dynamics (Haddad and Muir, 2021; Du, Hébert, and Huber, 2023).

A distinguishing feature of our paper is the use of traded risk factors to inform demand propagation across 17 currencies and 5 major non-FX asset classes. The mechanism rests on three empirically measurable objects: how demand shocks change non-diversifiable risks (q); how prices adjust to compensate intermediaries for absorbing additional risks (λ); and how each asset is exposed to these risks (β). Integrating these elements reveals rich transmission patterns, where demand shocks in one market propagate to others with varying magnitudes and even directions. As asset markets become increasingly interconnected, understanding how demand shocks propagate through common risk exposure is crucial for predicting and managing systematic market dynamics.

Bibliography

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

An, Y., and Z. Zheng. 2025. Information from uninformed demand: How predictable demand propagates across assets. Working Paper.

Andrews, I., J. H. Stock, and L. Sun. 2019. Weak instruments in IV regression: Theory and practice. *Annual Review of Economics* 11:727–53.

Balasubramaniam, V., J. Y. Campbell, T. Ramadorai, and B. Ranish. 2023. Who owns what? a factor model for direct stockholding. *Journal of Finance* 78:1545–91.

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

Ben-David, I., J. Li, A. Rossi, and Y. Song. 2022. Ratings-driven demand and systematic price fluctuations. *Review of Financial Studies* 35:2790–838.

BIS. 2022. Triennial central bank survey: OTC foreign exchange turnover in April 2022.

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

Brunnermeier, M. K., and Y. Sannikov. 2014. A macroeconomic model with a financial sector. *American Economic Review* 104:379–421.

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

Campbell, J. Y., and A. S. Kyle. 1993. Smart money, noise trading and stock price behaviour. *The Review of Economic Studies* 60:1–34.

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

Chaudhry, A., and J. Li. 2025. Endogenous elasticities: Price multipliers are smaller for larger demand shocks. Working Paper.

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

Cochrane, J. H. 2011. Presidential address: Discount rates. *The Journal of Finance* 66:1047–108.

Constantinides, G. M., J. C. Jackwerth, and A. Savov. 2013. The puzzle of index option returns. *The Review of Asset Pricing Studies* 3:229–57.

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.

Duffie, D. 2010. Asset price dynamics with slow-moving capital. *Journal of Finance* 65:1238–68.

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.

Fuchs, W., S. Fukuda, and D. Neuhann. 2025. Demand-system asset pricing: Theoretical foundations. *Available at SSRN 4672473* .

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.

Hartzmark, S. M., and D. H. Solomon. 2024. Marketwide predictable price pressure. Working Paper.

Hasbrouck, J. 1991. Measuring the information content of stock trades. *Journal of Finance* 46:179–207.

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.

He, Z., B. Kelly, and A. Manela. 2017. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics* 126:1–35.

He, Z., and A. Krishnamurthy. 2013. Intermediary asset pricing. *American Economic Review* 103:732–70.

———. 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.

Kozak, S., S. Nagel, and S. Santosh. 2018. Interpreting factor models. *The Journal of Finance* 73:1183–223.

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.

Li, J., and Z. Lin. 2022. Price multipliers are larger at more aggregate levels. Working Paper.

Liao, G. Y., and T. Zhang. 2025. The hedging channel of exchange rate determination. *The Review of Financial Studies* 38:1–38.

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.

Lo, A. W., and J. Wang. 2000. Trading volume: definitions, data analysis, and implications of portfolio theory. *Review of Financial Studies* 13:257–300.

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.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf. 2012. Currency momentum strategies. *Journal of Financial Economics* 106:660–84.

Milgrom, P., and N. Stokey. 1982. Information, trade and common knowledge. *Journal of Economic Theory* 26:17–27.

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.

Rostek, M. J., and J. H. Yoon. 2023. Imperfect competition in financial markets: Recent developments. Working Paper.

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

Shanken, J. 1992. On the estimation of beta-pricing models. *Review of Financial Studies* 5:1–33.

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. 2021. Price multipliers of anticipated and unanticipated shocks to demand and supply. Discussion paper.

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

Wachtel, P., and J. Young. 1990. The impact of treasury auction announcements on interest rates. *Quarterly Review of Economics and Business* 30.

Wang, J. 1993. A model of intertemporal asset prices under asymmetric information. *The Review of Economic Studies* 60:249–82.

A Appendix for Proofs

This appendix provides proofs omitted in the main text.

A.1 Proof of Proposition 1

Because factors have uncorrelated returns, we can project the return of any currency n onto the factors and obtain

$$r_n = \sum_{k=1}^{K} \beta_{n,k} r_k^{\text{factor}} + e_n, \tag{A1}$$

where e_n is the idiosyncratic return of currency n that is uncorrelated with any factors. Hence, by the law of one price and equation (11), the price impact of currency n is

$$\Delta p_n = \sum_{k=1}^{K} \beta_{n,k} \Delta p_k^{\text{factor}} = \sum_{k=1}^{K} \lambda_k \Delta Q_k^{\text{factor}} \text{var}(r_k^{\text{factor}}) \beta_{n,k}. \tag{A2}$$

Therefore, we have

$$\frac{\partial \Delta p_n}{\partial \Delta Q_k^{\text{factor}}} = \frac{\partial \Delta p_k^{\text{factor}}}{\partial \Delta Q_k^{\text{factor}}} \times \frac{\partial \Delta p_n}{\partial \Delta p_k^{\text{factor}}} = \lambda_k \text{var}(r_k^{\text{factor}}) \times \beta_{n,k}. \tag{A3}$$

Next, equation (7) implies that $\partial \Delta Q_k^{\text{factor}} / \partial \Delta Q_m = \beta_{m,k}$. Hence, we have proved

$$\frac{\partial \Delta p_n}{\partial \Delta Q_m} = \sum_{k=1}^{K} \frac{\partial \Delta Q_k^{\text{factor}}}{\partial \Delta Q_m} \times \frac{\partial \Delta p_n}{\partial \Delta Q_k^{\text{factor}}} = \sum_{k=1}^{K} \beta_{m,k} \times \lambda_k \text{var}(r_k^{\text{factor}}) \times \beta_{n,k}. \tag{A4}$$

A.2 Proof of Proposition 2

Let the return vector be $\mathbf{r} = (r_1, r_2, \ldots, r_N)^\top$ and the vector of currency flows $\Delta \mathbf{Q} = (q_1, q_2, \ldots, q_N)^\top$.

First, take the spectral decomposition $\text{var}(\mathbf{r}) = \mathbf{V} \mathbf{D} \mathbf{V}^\top$, where $\mathbf{D}$ is an $L \times L$ diagonal matrix with $L = \text{rank}(\text{var}(\mathbf{r}))$, and $\mathbf{V}$ is an $N \times L$ orthonormal matrix ($\mathbf{V}^\top \mathbf{V} = \mathbf{I}_L$). Next, perform another spectral decomposition:

$$\mathbf{D}^{1/2} \mathbf{V}^\top \text{var}(\Delta \mathbf{Q}) \mathbf{V} \mathbf{D}^{1/2} \mathbf{G} = \mathbf{G} \boldsymbol{\Pi}, \tag{A5}$$

where $\boldsymbol{\Pi}$ is a $K \times K$ diagonal matrix with entries ordered from largest to smallest, and $\mathbf{G}$ is a $K \times K$ orthogonal matrix ($\mathbf{G}^\top \mathbf{G} = \mathbf{I}_K$). Define the portfolio-weight matrix $\mathbf{W} = \mathbf{V} \mathbf{D}^{-1/2} \mathbf{G}$.

We claim that the factors constructed by the portfolio weight matrix $\mathbf{W}$ satisfy the

conditions (17) and (18). First,

$$\text{var}(\mathbf{W}^\top \mathbf{r}) = \mathbf{G}^\top \mathbf{D}^{-\frac{1}{2}} \mathbf{V}^\top \mathbf{V} \mathbf{D} \mathbf{V}^\top \mathbf{V} \mathbf{D}^{-\frac{1}{2}} \mathbf{G} = \mathbf{I}_K. \tag{A6}$$

Thus each factor return is uncorrelated with the others and has unit variance.

Second, equation (7) in matrix form yields

$$\Delta \mathbf{Q}^{\text{factor}} = \text{var}(\mathbf{W}^\top \mathbf{r})^{-1} \text{cov}(\mathbf{W}^\top \mathbf{r}, \Delta \mathbf{Q}^\top \mathbf{r}) = \mathbf{G}^\top \mathbf{D}^{\frac{1}{2}} \mathbf{V}^\top \Delta \mathbf{Q}, \tag{A7}$$

where we use $\mathbf{W}^\top \mathbf{r} = \mathbf{I}_K$, $\mathbf{W} = \mathbf{V} \mathbf{D}^{-\frac{1}{2}} \mathbf{G}$, and $\text{var}(\mathbf{r}) = \mathbf{V} \mathbf{D} \mathbf{V}^\top$. Hence,

$$\text{var}(\Delta \mathbf{Q}^{\text{factor}}) = \mathbf{G}^\top \mathbf{D}^{\frac{1}{2}} \mathbf{V}^\top \text{var}(\Delta \mathbf{Q}) \mathbf{V} \mathbf{D}^{\frac{1}{2}} \mathbf{G} = \mathbf{\Pi}, \tag{A8}$$

which is diagonal; factor flows are therefore uncorrelated, and each diagonal element of $\mathbf{\Pi}$ is its variance.

The factors can be scaled arbitrarily, but the product $\text{var}(\Delta \mathbf{Q}^{\text{factor}})\, \text{var}(\mathbf{W}^\top \mathbf{r})$ is invariant, because doubling a factor's weights doubles its return while halving its flow. We adopt the convenient normalization $\text{var}(\mathbf{W}^\top \mathbf{r}) = \mathbf{I}_K$; ordering the factors by the diagonal elements of $\text{var}(\Delta \mathbf{Q}^{\text{factor}}) = \mathbf{\Pi}$ therefore guarantees that the leading factors explain the largest share of trading-induced risk.

A.3 Invariance of Factors under Alternative Numeraire Currency

This appendix shows that the factors constructed in Appendix A.2 are invariant to the numeraire currency used to measure returns and flows.

Assume we switch the numeraire from USD to currency N. Let $\Delta \tilde{Q}_n$ be the flow from currency N into currency n for $n = 1, \ldots, N-1$, and let $\Delta \tilde{Q}_N$ be the shock from currency N into USD. Recall that ΔQ_n denotes the shock from USD into currency n. Because each USD-based trade decomposes into a leg from USD to currency N and a leg from currency N to the target currency, the shocks transform according to

$$\Delta \tilde{\mathbf{Q}} = (\Delta \tilde{Q}_1, \Delta \tilde{Q}_2, \ldots, \Delta \tilde{Q}_{N-1}, \Delta \tilde{Q}_N)^\top = \left(\Delta Q_1, \Delta Q_2, \ldots, \Delta Q_{N-1}, -\sum_{n=1}^{N} \Delta Q_n \right)^\top = \mathbf{C} \Delta \mathbf{Q}, \tag{A9}$$

where we define the matrix

$$
\mathbf{C} := \begin{pmatrix}
1 & 0 & \ldots & 0 & 0 \\
0 & 1 & \ldots & 0 & 0 \\
\vdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & \ldots & 1 & 0 \\
-1 & -1 & \ldots & -1 & -1
\end{pmatrix}. \tag{A10}
$$

Returns transform analogously. For $n = 1, \ldots, N-1$, define $\tilde{r}_n$ as the excess return from borrowing at currency N's risk-free rate and investing at currency n's risk-free rate; let $\tilde{r}_N$ be the corresponding return on USD funding. Then,

$$
\tilde{\mathbf{r}} = (\tilde{r}_1, \tilde{r}_2, \ldots, \tilde{r}_{N-1}, \tilde{r}_N)^\top = (r_1 - r_N, r_2 - r_N, \ldots, r_{N-1} - r_N, -r_N)^\top = \mathbf{C}^\top \mathbf{r}. \tag{A11}
$$

A crucial property of $\mathbf{C}$ is $\mathbf{CC} = \mathbf{I}_N$, which follows directly from the definition above. Intuitively, changing the numeraire from USD to currency N and then back again leaves all quantities unchanged, so multiplying $\mathbf{C}$ by itself must be identity.

Next, we apply the procedure in Appendix A.2 to the transformed data $\tilde{\mathbf{r}}$ and $\Delta\tilde{\mathbf{Q}}$ and show that it produces the same factors obtained from $\mathbf{r}$ and $\Delta\mathbf{Q}$. Let the spectral decomposition be $\mathrm{var}(\tilde{\mathbf{r}}) = \tilde{\mathbf{V}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^\top$ and $\mathrm{var}(\mathbf{r}) = \mathbf{V}\mathbf{D}\mathbf{V}^\top$. From (A11) we also have $\mathrm{var}(\tilde{\mathbf{r}}) = \mathbf{C}^\top \mathrm{var}(\mathbf{r})\mathbf{C}$. Hence there exists an $L \times L$ orthogonal matrix $\mathbf{O}$ ($\mathbf{O}^\top\mathbf{O} = \mathbf{I}_L$) such that

$$
\mathbf{C}^\top \mathbf{V}\mathbf{D}^{\frac{1}{2}} = \tilde{\mathbf{V}}\tilde{\mathbf{D}}^{\frac{1}{2}}\mathbf{O}. \tag{A12}
$$

Using (A9) and the fact that $\mathbf{CC} = \mathbf{I}_N$,

$$
\tilde{\mathbf{D}}^{1/2}\tilde{\mathbf{V}}^\top \mathrm{var}(\Delta\tilde{\mathbf{Q}})\, \tilde{\mathbf{V}}\tilde{\mathbf{D}}^{1/2} = \mathbf{O}\mathbf{D}^{\frac{1}{2}}\mathbf{V}^\top \mathbf{CC}\, \mathrm{var}(\Delta\mathbf{Q})\, \mathbf{C}^\top\mathbf{C}^\top\mathbf{V}\mathbf{D}^{\frac{1}{2}}\mathbf{O}^\top = \mathbf{O}\mathbf{D}^{\frac{1}{2}}\mathbf{V}^\top \mathrm{var}(\Delta\mathbf{Q})\, \mathbf{V}\mathbf{D}^{\frac{1}{2}}\mathbf{O}^\top. \tag{A13}
$$

The spectral decomposition in equation (A5) therefore read

$$
\tilde{\mathbf{D}}^{1/2}\tilde{\mathbf{V}}^\top \mathrm{var}(\Delta\tilde{\mathbf{Q}})\, \tilde{\mathbf{V}}\tilde{\mathbf{D}}^{1/2}\tilde{\mathbf{G}} = \tilde{\mathbf{G}}\tilde{\mathbf{\Pi}}, \tag{A14}
$$

$$
\mathbf{D}^{\frac{1}{2}}\mathbf{V}^\top \mathrm{var}(\Delta\mathbf{Q})\, \mathbf{V}\mathbf{D}^{\frac{1}{2}}\mathbf{G} = \mathbf{G}\mathbf{\Pi}, \tag{A15}
$$

which implies $\mathbf{G} = \mathbf{O}^\top\tilde{\mathbf{G}}$ and $\tilde{\mathbf{\Pi}} = \mathbf{\Pi}$. Thus, the eigenvalues are identical and the eigenvectors differ only by an orthogonal rotation.

Hence the factor–portfolio weights under the new numeraire are $\tilde{\mathbf{W}} = \tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-1/2}\tilde{\mathbf{G}}$, and

the corresponding factor returns satisfy

$$\tilde{\mathbf{W}}^\top \tilde{\mathbf{r}} = (\mathbf{C}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-\frac{1}{2}}\tilde{\mathbf{G}})^\top \mathbf{r} = (\mathbf{C}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-\frac{1}{2}}\mathbf{O}\mathbf{G})^\top \mathbf{r} = (\mathbf{V}^\top \mathbf{C}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-\frac{1}{2}}\mathbf{O}\mathbf{G})^\top \mathbf{f}, \qquad \text{(A16)}$$

where the last step uses $\mathbf{r} = \mathbf{V}\mathbf{f}$ for some L-dimensional vector $\mathbf{f}$, because $\mathrm{rank}(\mathrm{var}(\mathbf{r})) = L$.

From equation (A12), we we left-multiply both sides by $\tilde{\mathbf{D}}^{-\frac{1}{2}}\tilde{\mathbf{V}}^\top$ to obtain $\tilde{\mathbf{D}}^{-\frac{1}{2}}\tilde{\mathbf{V}}^\top \mathbf{C}^\top \mathbf{V}\mathbf{D}^{\frac{1}{2}} = \mathbf{O}$ (recall that $\tilde{\mathbf{V}}^\top \tilde{\mathbf{V}} = \mathbf{I}_L$). Taking transposes gives $\mathbf{D}^{\frac{1}{2}}\mathbf{V}^\top \mathbf{C}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-\frac{1}{2}} = \mathbf{O}^\top$. Left-multiplying by $\mathbf{D}^{-1/2}$ and right-multiplying by $\mathbf{O}$ yields $\mathbf{V}^\top \mathbf{C}\tilde{\mathbf{V}}\tilde{\mathbf{D}}^{-\frac{1}{2}}\mathbf{O} = \mathbf{D}^{-\frac{1}{2}}$ (recall that $\mathbf{O}^\top \mathbf{O} = \mathbf{I}_L$).

Substituting this identity into (A16) gives

$$\tilde{\mathbf{W}}^\top \tilde{\mathbf{r}} = (\mathbf{D}^{-\frac{1}{2}}\mathbf{G})^\top \mathbf{f} = (\mathbf{V}\mathbf{D}^{-1/2}\mathbf{G})^\top \mathbf{V}\mathbf{f} = \mathbf{W}^\top \mathbf{r}, \qquad \text{(A17)}$$

where we used $\mathbf{V}^\top \mathbf{V} = \mathbf{I}_L$. Therefore the factor returns—and hence the factors themselves—are identical under either numeraire.

# B Fraction of Return Variance Explained by Factor Flows

This appendix measures how much of each factor's return variance comes from factor flows. Equilibrium flows may carry information, so price impacts have two parts: a permanent information component and a temporary demand component. Our estimate λ_k captures the demand component, and we have shown that its impact reverts within one month. We now compute, across different horizons, the share of return variance explained by this temporary effect. As the horizon lengthens, the temporary impact decays and accounts for a smaller fraction of total variance.

Consider stochastic trading flow over the horizon $[0, T]$. During an infinitesimal interval dt, the flow shock for factor k has standard deviation $\sigma(\Delta Q_{k,t}^{\text{factor}})\sqrt{dt}$, where $\sigma(\Delta Q_{k,t}^{\text{factor}})$ is the annualized flow volatility. Each unit of flow moves the price by $\lambda_k \sigma^2(r_{k,t}^{\text{factor}})$, with $\sigma(r_{k,t}^{\text{factor}})$ denoting the annualized return volatility. Hence the contemporaneous price change generated by the shock equals $\lambda_k \, \sigma^2(r_{k,t}^{\text{factor}}) \, \sigma(\Delta Q_{k,t}^{\text{factor}})\sqrt{dt}$. Assuming this impact decays linearly to zero within one month (1/12 years), a shock arriving at time t still affects the price at the terminal date T by the factor $\max\{0, 1-12(T-t)\}$. Accordingly, the proportion of the terminal price variance attributable to order flow is

$$\frac{\int_0^T \big(\max\{0, \, 1 - 12(T - t)\}\lambda_k \sigma^2(r_{k,t}^{\text{factor}})\sigma(\Delta Q_{k,t}^{\text{factor}})\big)^2 \, dt}{\int_0^T \sigma^2(r_{k,t}^{\text{factor}})dt}, \qquad \text{(A18)}$$

which measures the share of total variance over $[0, T]$ that arises from temporary price

Figure A1: **Factor flow contribution to return variance**

impacts of order-flow shocks. This fraction simplifies to

$$
\begin{cases}
\lambda_k^2 \sigma^2(r_{k,t}^{\text{factor}}) \sigma^2(\Delta Q_{k,t}^{\text{factor}})(1 - 12T + 48T^2) & 0 < T \leq \dfrac{1}{12}, \\[2ex]
\dfrac{\lambda_k^2 \sigma^2(r_{k,t}^{\text{factor}}) \sigma^2(\Delta Q_{k,t}^{\text{factor}})}{36T} & T > \dfrac{1}{12}.
\end{cases}
\tag{A19}
$$

Figure A1 shows that for the three factors, flows explain roughly 10–35% of return variance at the 1-week horizon, 5–15% at one month, and rapidly less at longer horizons.

Supplemental Appendix of
"Demand Propagation Through Traded Risk Factors"

A   Inclusion of Non-spot FX Derivatives Trading Flows

Foreign exchange trades can be executed in the spot market and in the derivatives market of forwards and swaps. Trading in the derivatives market can expose the intermediary to foreign exchange risk. Consider a customer-initiated trade of selling $100-worth of JPY 1-month forward against USD. In the absence of other trades, an intermediary who has no capital, maintains a net neutral FX exposure, and serves as the counterparty in this trade, must satisfy the obligation to deliver $100 in a month by setting aside $100/(1 + r_{1M}^{\$})$ today, where $r_{1M}^{\$}$ is the 1-month USD risk-free rate. Similarly, the intermediary will sell $100/(1 + r_{1M}^{\text{JPY}})$ of JPY today to both fund his USD purchase and to ensure FX neutrality when he receives the promised delivery from the customer. To the intermediary, therefore, a forward contract is no different from a spot transaction but for the fact that the amount of implied FX exposure in a forward is less than its notional.

Because we are interested in measuring all the FX risks that intermediaries have to bear by accommodating customer trading flows, we need to consider trading flows in both the spot and the derivatives market.[38] In this appendix, we explore the difference between trading flows into the spot versus the derivatives market and the implications of using trading data in only one of the two markets in our analysis.

We start by examining the observed trading flows into individual currencies. The triennial survey conducted by the Bank of International Settlement (BIS) indicates that there is twice as much trading flow in the FX derivatives market as in the spot market (Appendix Figure SA1). Appendix Table SA1 reports the correlation between the net flow into the spot versus the derivatives market for each of the six major currencies in our sample. The absolute strength of the correlation ranges between 0.17 and 0.62, suggesting sizeable comovements in trading flows between the spot and the derivatives FX market.

Comovements in observed trading flows could be induced by common risk factors that are present in both the spot and the derivatives market. If so, trading data from either market alone should be sufficient to recover the traded FX risk factors. At the same time, if there are strong comovements in trading flows to the traded FX factors, then relying on data from only one market risks introducing measurement error in the estimation of price sensitivity to risks.

---

[38]We treat swaps as a spot transaction plus a forward contract.

Figure SA1: **FX Daily Turnover Against USD**

Notes: This figure plots the global daily volume of foreign exchange spot versus forward and FX swaps 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 BIS.

Table SA1: **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.

In Appendix Table SA2, we compare the traded FX factors recovered separately from the spot market and the non-spot derivatives market. The top row shows the correlation between *returns* of factors estimated using only one of the individual markets. For the first factor, the return correlation is close to 1, and this correlation is 77% for the second factor and 73% for the third factor. Such pronounced relationships underscore the robustness of the underlying factors and suggest that the same risk factors drive trading across the spot and the derivatives market. The bottom row shows the correlation between *flows* to factors estimated using only one of the individual markets. The correlations are -0.51, -0.13, and -0.35 for the three factors, respectively.

The marked association between factor returns and factor flows points to the strength and limitation of using only data in the spot market. On the one hand, the tight correlation

Table SA2: **Correlation between Returns and Flows to Factors Estimated in Different Samples**

	Factor 1	Factor 2	Factor 3
Return	0.99	0.77	0.73
Flow	-0.51	-0.13	-0.35

Notes: This table reports the correlation between the returns and flows to each of the top three traded risk factors as estimated in the spot market versus in the non-spot derivatives market.

between factor returns constructed using data from individual markets shows that the spot market alone is sufficient to recover the underlying risk factors because these factors drive trading in both the spot and derivatives markets. On the other hand, using only data from the spot market is likely insufficient for estimating these factors' price sensitivity to risks because the spot market data alone may not provide an appropriate measure of the flow changes. Estimating price sensitivity to risks requires instrumenting for the flow that induces the observed price change. As spot flows and derivatives flows are highly correlated, it is empirically difficult to isolate variations in just the spot flow. Specifically, because factor flows in the spot market are negatively correlated with factor flows in the derivatives market, instrumenting for just the spot market will overestimate factor flows, biasing the estimate to imply smaller price sensitivity to risks.

B   Additional Figures and Tables

Table SA3: **Correlation Between Traded FX 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 FX factors constructed based on the full sample versus returns and flows of the traded FX 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.

Table SA4: **Demand Propagation Between Traded FX Factors and Non-FX Asset Classes**

	CDS	Comm	CorpBond	Opt	UST
Dollar	-2.0	-5.0	-2.8	-4.4	-0.5
Carry	3.7	1.6	3.7	6.1	-2.3
Euro-Yen Residual	-2.5	-10.3	-7.3	-6.6	-1.6

Notes: This table uses Proposition 1, the estimated factor-level price sensitivity to risks λ_k from Table 5, and the beta loadings of assets to factors (signs illustrated in Figure 3) to compute cross-multiplier between traded FX factors and six non-FX asset classes. Each entry represents the price movement in bps of a column asset, as induced by a $1 billion demand shock into a traded FX factor.

Figure SA2: **Cumulative Flow by Investor Type to Traded FX Factors**
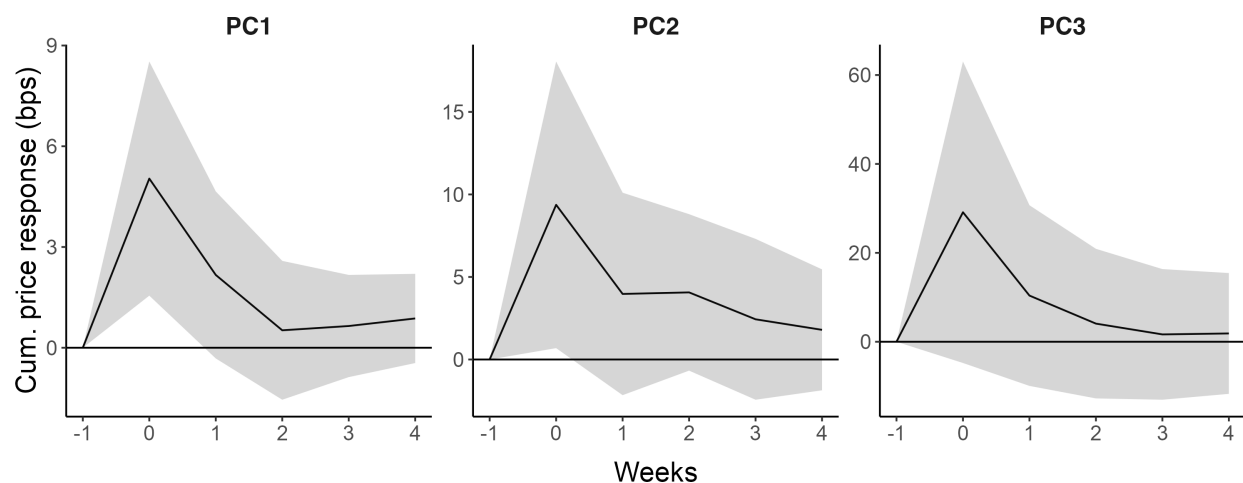

(a) **Dollar Factor**

(b) **Carry Factor**

(c) **Euro-Yen Factor**

Notes: This figure displays the cumulative flows by customer type into the top three traded FX factors between September 2012 and December 2023. The Net Total represents the net customer flows that Banks (intermediaries) need to absorb.

SA.5

Figure SA3: **Reversion of Contemporaneous Price Response**

Notes: This figure shows the cumulative price responses for the traded FX factors. These responses, measured per billion of demand shocks, are estimated by regressing the return from week $t-1$ to $t+h$ (for $h = 0, 1, 2, 3, 4$) on the instrumented flow from week $t-1$ to t. 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.

Figure SA4: **Decomposition of Asset-Class Returns Explained by Traded FX Factors Outside of Crises**

Notes: This figure decomposes the returns of individual assets into the Dollar, the Carry, and the Euro-Yen Residual factors. The decomposition is achieved by regressing asset class monthly average excess return between 2000-02 and 2023-12 on returns from the three traded FX factors. We exclude the GFC (2007-07 through 2010-07) and COVID (2020-01 through 2020-06) period. The returns from CDS are available starting 2007-04. The returns from Opt end in 2022-12. It reports both the marginal R^2 values attributed to each factor and the total R^2. The positive and negative signs illustrate the direction of the beta loadings.

SA.6