Intermediary Elasticity and Limited Risk-Bearing Capacity*

Yu An[†] Amy W. Huber[‡]

This version: October 2024
Download most recent version here.

Abstract

We study intermediaries' limited risk-bearing capacity and its implications for asset prices. We introduce a new measure, "intermediary elasticity", defined as the price sensitivity to a marginal increase in risk induced by trading demand shocks. We apply our framework to the foreign exchange (FX) market and find that just three traded risk factors can jointly account for 90% of the non-diversifiable risks borne by intermediaries when accommodating FX trading flows. These three traded risk factors resemble the Dollar, the Carry, and the Euro-Yen, and reveal that intermediaries accumulated $0.8 trillion in exposure to the Carry over the last decade. Through instrumental variable analysis, we show that intermediaries raise prices by 5 to 30 bps in response to a $1 billion trading demand shock to these factors. We use our estimated FX-factor elasticity to quantify the cross-elasticity of a panel of currencies and across 7 major asset classes.

JEL Classifications: G11, G12, G15, G2
Keywords: Elasticity, Risk, Intermediary, FX, Traded Factor

*We thank Itamar Drechsler, Thomas Gruenthaler, Ben Hébert, Arvind Krishnamurthy, Jiacui Li, Matteo Maggiori, Angelo Ranaldo, Nikolai Roussanov, Davide Tomio, Adrien Verdelhan, Moto Yogo, Tony Zhang, and seminar participants at Temple University, JHU Carey, Wharton, SAFE asset pricing workshop, Princeton, WAPFIN, Fed Board, WAIFS for comments. We gratefully acknowledge funding from the Wharton Dean's Research Fund, the Jacobs Levy Equity Management Center for Quantitative Financial Research at the Wharton School of the University of Pennsylvania, and the General Research Support Fund at Johns Hopkins Carey Business School. All remaining errors are our own.

†Johns Hopkins Carey Business School. Email: yua@jhu.edu.
‡The Wharton School of the University of Pennsylvania. Email: amyhuber@wharton.upenn.edu.

1 Introduction

A growing literature offers evidence that financial intermediaries are central to asset pricing (e.g., Haddad and Muir, 2021; Du, Hébert, and Huber, 2022). Intermediaries can have limited risk-bearing capacity due to reasons like risk aversion or liquidity constraints (e.g., Gabaix and Maggiori, 2015; Kondor and Vayanos, 2019). Hence, in contrast to the traditional asset pricing paradigm, even trading demand shocks that are not motivated by or informative of fundamentals, can affect asset prices when such shocks are absorbed by intermediaries (e.g., Froot and Ramadorai, 2008). In this paper, we study intermediaries' risk-bearing capacity in the foreign exchange (FX) market and quantify trading's impact on prices across currencies and asset classes. Our results underscore the role that intermediaries play in driving cross-asset pricing dynamics.

Central to our investigation is the concept of "intermediary inverse semi-elasticity", which we define as the price sensitivity to a marginal increase in risk induced by trading demand shocks. For simplicity, we abbreviate it as "intermediary elasticity."[1] Our concept builds upon the classic asset pricing framework, where asset prices are driven by risks. We specialize this framework to trading-induced risks to empirically measure the *changes* in asset prices due to *changes* in risks. Relative to other measures of elasticity in asset pricing (e.g., Koijen and Yogo, 2019; Huber, 2023), "intermediary elasticity" is defined directly with respect to *risks*, which makes it a natural measure of risk-bearing capacity. If intermediary elasticity is zero, then intermediaries are able to perfectly absorb risks from accommodating trading demand

[1]Elasticity measures how responsive one variable is to changes in another variable, and is typically expressed as a percentage change in one variable relative to the percentage change in another. Semi-elasticity measures change in the independent variable (denominator) in amount and not in percentages. Inverse elasticity refers to when price is the dependent variable (numerator). "Price multiplier" is another term that is sometimes used to refer to inverse elasticity, e.g., in Gabaix and Koijen (2021).

shocks and the risk-bearing capacity is unlimited. Conversely, the larger the intermediary elasticity, the more limited the risk-bearing capacity.

Intermediary elasticity is a particularly relevant concept in the FX market. Almost all FX trades are intermediated by banks, dealers, or hedge funds (Chaboud, Rime, and Sushko, 2023). As FX intermediaries are sophisticated agents who actively trade among themselves to share risks, currency prices exhibit a strong factor structure (Lustig and Verdelhan, 2007), echoing the classic notion that only non-diversifiable risks are priced (Markowitz, 1952; Ross, 1976). In such a market, should trading demand shocks affect asset prices, they do so by altering the amount of non-diversifiable risks that intermediaries must bear. The key question is: what priced risks are affected by trading? We develop a novel technique that extends the classic arbitrage pricing theory to empirically identify risk factors with the most trading-induced risks, and we estimate the intermediary elasticity of these *traded* risk factors.

We employ a unique dataset of global intermediaries' FX trading to study these intermediaries' risk-bearing capacity. We obtain daily trades between intermediaries and their customers in not just FX spot but also FX forward and FX swap from the CLS Group, the largest single source of FX executed data available to the market. We first identify the three most important traded FX risk factors ("traded FX factors"). These traded FX factors jointly account for 90% of the non-diversifiable risks that intermediaries bear when accommodating FX trading flows and they resemble the Dollar, the Carry, and the Euro-Yen.[2] We then use instrumental variables to estimate the intermediary elasticity of these three traded FX factors, and find that FX intermediaries have rather limited risk-bearing capacity. Our factor-specific intermediary elasticity allows us to recover the cross-elasticity

[2]The Dollar factor goes long in U.S. dollar and shorts all other currencies. The Carry factor goes long in high interest rate currencies such as the Australian dollar and shorts low interest rate currencies such as the euro (EUR) and the Japanese yen (JPY). The Euro-Yen factor goes long in EUR and shorts JPY, while being neutral on all other currencies.

between a whole panel of currencies, a first in the literature. Finally, we show that returns in CDS, commodities, corporate bonds, equities, options, and U.S. Treasury bonds can also be explained by the three traded FX factors, and we apply the estimated FX-factor elasticity to obtain novel estimates of cross-elasticity between these asset classes.

We start by identifying FX risk factors that have the most trading-induced risks in 16 non-U.S. dollar (USD) currencies. We employ a modified principal component analysis (PCA) procedure on a weekly panel of trading flows and returns to account for variations in flows and returns jointly. Our procedure contrasts with the standard PCA that rationalizes the covariance structure of either flows or returns in isolation. A standard PCA on FX trading flows alone, for example, simply points to portfolios with the most traded currencies because it neglects common variations in currency returns. A standard PCA on FX return alone, on the other hand, can surface unconditional risk factors priced in FX, but is silent on whether these factors are traded.[3] Remarkably, we find that the three most important traded FX factors resemble the Dollar, the Carry, and the Euro-Yen, where the first two are also the most important unconditional risk factors.[4] As these factors capture non-diversifiable risks in FX trading, net flows into these factors measure the intermediaries' otherwise unobserved risk exposure from FX trading. For example, we find that intermediaries accumulated $0.8 trillion in exposure to the Carry between 2012 and 2023.

Having identified the traded FX factors, we proceed to estimate the intermediary elasticity of these factors. By construction, our traded FX factors have uncorrelated returns, meaning that intermediaries view these factors as uncorrelated risks. We thus estimate the

[3]In fact, the risk premium of these unconditional risk factors is typically microfounded on consumption-based models with stochastic discount factors (SDFs) that do not rely on investor trading (Lustig and Verdelhan, 2007).

[4]The Euro-Yen is a new factor that we propose. This factor is important because there is active trading between EUR and JPY (either directly or intermediated through USD), but the two currencies are traded in the same direction in both the Dollar and the Carry factors.

intermediary elasticity factor-by-factor without worrying about cross-factor substitution. We must, however, instrument for trading demand shocks in each of the factors to purge the effect of fundamentals (e.g., the arrival of new information) on the observed price. We employ as instrumental variables the announcements of the offering amount at upcoming sovereign bond auctions in the U.S., Australia, Canada, the U.K., Japan, Italy, France, and Germany. These sovereign auctions often attract foreign investors who need to convert currencies to participate, making the instruments relevant. At the same time, because these auctions are typically forward-guided, there is likely limited information in the *announcement* of the auction, making the instruments plausibly exogenous and meet the exclusion restriction.[5] We estimate that to induce intermediaries to absorb a $1 billion trading demand shock, the price needs to rise by 5 basis points (bps) for the Dollar, 9 bps for the Carry, and 29 bps for the Euro-Yen. The price response to shocks in the traded FX factors is much larger than the estimated response to trading demand shocks in the U.S. equities market factor.[6] Our result highlights the rather limited risk-bearing capacity in FX. Viewed through the lens of our model, the high intermediary elasticity in FX could reflect limited arbitrage capital, leading to large price responses when intermediaries absorb non-diversifiable risks.[7] In fact,

[5]First, the instruments would not be exogenous if they contained news about fundamentals that moves exchange rates. Such news, if present, likely also moves bond yields. Empirically, Wachtel and Young (1990) find that while the results from Treasury bond auctions significantly affect bond yields, the announcements of these auctions have no detectable effect on yields. Second, the instruments would not meet the exclusion restriction if exchange rates reacted to trading because intermediaries learned of customers' private information, rather than because intermediaries have limited risk-bearing capacity. Such private information, if present, is likely revealed in customers' demand of the auctioned bonds, which is not known at announcements.

[6]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 and its effect is thus comparable to intermediary elasticity. The average market capitalization between 2012 and 2022 is about $31.7 trillion. A $1 billion trading demand shock in our sample period therefore raises the price of the market factor by about 2 bps.

[7]Although the trading volume in FX is enormous, 75% of the trading is conducted among intermediaries, according to the Bank for International Settlements's Triennial Central Bank Survey (BIS, 2022).

the cross-factor variation in the estimated intermediary elasticity could reflect differences in available arbitrage capital across risk factors, whereby the lesser-known factors such as the Euro-Yen have more limited arbitrage capital.

The estimated intermediary elasticity allows us to compute the (inverse semi) cross-elasticity between any pair of currencies. We use "cross-elasticity" to measure the impact of a trading demand shock to currency A on the price of currency B, while holding the trading demand shocks of all other currencies constant. Estimating such cross-elasticity for a whole panel of currencies is challenging, as currencies can be seen as substitutable in an investment portfolio and likely have correlated trading demand shocks. Our insight lies in mapping the cross-elasticity of currencies to currencies' exposure to traded FX factors. When intermediaries accommodate trading demand shocks to currency A, they bear additional non-diversifiable factor-level risks. These risks influence factor prices through the estimated intermediary elasticity and ultimately affect the price of currency B via the law of one price. As factor loadings differ both in direction and in magnitude across currencies, we recover currencies' rich cross-substitution pattern. We find large cross-elasticity between the Australian dollar (AUD) and the Canadian dollar (CAD) because these two currencies are traded in the same direction in all three traded FX factors. In contrast, the cross-elasticity between the Japanese yen (JPY) and either AUD or CAD is small because JPY and these two currencies are on opposite sides of the Carry trade and hedge each other in exposure to the Carry factor. Accordingly, although the euro (EUR) and JPY are both low interest-rate currencies, our estimates suggest that they are not the same: they "complement" each other in reducing the intermediary's exposure to the Euro-Yen factor and therefore have only modest cross-elasticity.

Finally, we use the traded FX factors to inform cross-elasticity between asset classes. We show that returns in six other asset classes load on the traded FX factors. Consequently, trading demand shocks originating in FX can affect prices in other asset classes through the traded FX factors. Similarly, a trading demand shock in, say, corporate bonds moves prices in corporate bonds and other assets through its effect on the traded FX factors. We find that trading demand shocks move the price the least in U.S. Treasury bonds (Treasurys), corroborating the observation that the Treasurys market is deep and liquid. The Treasurys also stand out as the only asset that has negative cross-elasticity with other assets, reminiscent of Treasurys' "safe haven" status. In the context of our exercise, Treasurys' safe haven property arises because it is the only asset that loads negatively on the Carry factor.

We caution that our estimates only capture the cross-elasticity channeled through the three traded FX factors. These factors explain about 80% of the unconditional return variation in FX currencies and about 30% of the unconditional return variation in non-FX assets. Our estimates thus miss potential cross-elasticity due to common exposure to other factors. Nevertheless, our estimates highlight that even though intermediaries are active in several markets, these markets do not necessarily move in tandem. Rather, understanding how asset markets are interconnected requires understanding which risk factors are affected by trading, how different assets are exposed to these traded risk factors, and what the intermediary elasticity of these factors is.

More generally, this paper advances the literature on intermediary-based asset pricing on two fronts. First, we present a novel way to measure the impact of intermediaries' risk-bearing capacity on asset prices. Risk aversion and various constraints such as equity capital, regulation, and liquidity could theoretically limit intermediaries' risk-bearing capacity (e.g., Gabaix and Maggiori, 2015; He and Krishnamurthy, 2017; Kondor and Vayanos, 2019). Empirical evidence shows that the risk of intermediaries' risk-bearing capacity tightening is priced (e.g., Du, Tepper, and Verdelhan, 2018; Du, Hébert, and Huber, 2022; Duffie, Fleming, Keane, Nelson, Shachar, and Van Tassel, 2023), but offers limited insight into how this risk-bearing capacity matters at the margin. Extending the portfolio theories of Markowitz (1952) and Ross (1976) to a representative intermediary (He and Krishnamurthy, 2013), we show that trading-induced non-diversifiable risks are priced by intermediaries at the margin. More importantly, we derive an original technique to identify directly from data the *traded* non-diversifiable risks, making it possible to quantify intermediaries' risk-bearing capacity by estimating the sensitivity of price change to changes in trading-induced risks ("intermediary elasticity"). Second, we illustrate the nuanced cross-asset pricing dynamics in intermediated financial markets. Intermediaries have been shown to matter for prices of many assets (e.g., Adrian, Etula, and Muir, 2014; He, Kelly, and Manela, 2017; Haddad and Muir, 2021). We use a factor-pricing framework to estimate the first set of cross-asset elasticity from the factor-level intermediary elasticity. Our results underscore that although trading demand shocks could propagate through assets and markets, such propagations need not be uniform. Understanding cross-asset dynamics through exposure to common factors is a framework that can be broadly applied. For example, tying the tax burden of financial transactions to risk exposure rather than to specific transactions is likely more effective in shaping agent's behaviors (Tobin, 1978).

Many findings in this paper relate directly to the literature on exchange rate determination. In particular, FX trading flows are shown to influence exchange rates above and beyond the information trading may convey (e.g., Evans and Lyons, 2002; Pasquariello, 2007; Froot and Ramadorai, 2008). We emphasize that flows matter because they push intermediaries against their risk-bearing capacity. Exchange rates have also been shown to affect and be affected by asset demands in equities and fixed income.[8] The distinguishing feature of our paper is that we let the data inform the specific risk factors that link FX trading with other asset markets. In fact, we find that the risk factors with the most trading-induced risks are also the ones that price the unconditional exchange rate returns (Lustig, Roussanov, and Verdelhan, 2011). By showing that these risk factors are traded by customers and priced by intermediaries at the margin, we uncover new facts about intermediaries' time-varying exposure to these factors, and complement the existing understanding of why these risks are priced into unconditional expected returns (e.g., Bansal and Dahlquist, 2000; Lustig and Verdelhan, 2007; Ready, Roussanov, and Ward, 2017).

Finally, this paper contributes to the growing literature that links trading demand with asset prices. Limited risk-bearing capacity has long been suggested (e.g., Grossman and Miller, 1988; Duffie, 2010; Vayanos and Vila, 2021) as an explanation for trading's observed impact on asset prices (e.g., Coval and Stafford, 2007; Lou, 2012). More recently, several papers provide reduced-form evidence in support of this channel.[9] Our key innovation is to use explicitly measured changes in the quantity of risk to pin down the relationship between price and risk at the margin. Our notion of "intermediary elasticity" builds on the burgeoning body of research that adapts concepts in industrial organization to understanding asset price dynamics (e.g., Koijen and Yogo, 2019; Bretscher, Schmid, Sen, and Sharma, 2022; Jiang, Richmond, and Zhang, 2024), but is uniquely defined with respect to risks as opposed to securities. Our approach is particularly appealing in understanding markets where agents are

[8] Camanho, Hau, and Rey (2022) examine the connection to the equity market through portfolio rebalancing. Liao and Zhang (2020), Jiang, Krishnamurthy, and Lustig (2021), and Gourinchas, Ray, and Vayanos (2024), among others, study the relationship with the bond markets due to hedging, safe asset demand, and preferred-habitat investors.

[9] Li and Lin (2022) document that price multipliers (inverse elasticity) are larger at more aggregate levels, where risks are more difficult to diversify away. Albuquerque, Cardoso-Costa, and Faias (2024) find that price elasticity correlates with return volatility. Wittwer and Allen (2024) find that price elasticity correlates with dealer capitalization.

free from institutional frictions (e.g., investment mandates) and can optimize over portfolios rather than individual securities. In so doing, we circumvent several limitations inherent in defining elasticity with respect to securities, including the need to introduce non-pecuniary preferences and to impose potentially rigid cross-substitution patterns.[10] Our methodology in part derives from An (2023) and is related to An, Su, and Wang (2024), who use a factor model with trading quantity to improve expected return estimation in the equity markets.

In the next section, we lay out our theoretical framework. We introduce the various sources of data we use in Section 3 and proceed to recover the traded FX factors in Section 4. We employ an instrumental variable approach to estimate the intermediary elasticity of the traded FX factors in Section 5, and apply these estimates to recover the cross-elasticity between currency pairs. We explore the connection between the FX market and six other asset classes in Section 6, and derive cross-elasticity between asset classes. We conclude in Section 7.

2 Theoretical Framework

This section first presents the conceptual framework of our study, and then details the construction of traded risk factors, the solution for intermediary elasticity, and the mapping from intermediary elasticity to cross-currency elasticity.

2.1 Model Setup and Conceptual Framework

There are three periods: $t = 0$, $t = 1$, and $t = 2$; and there are $N + 1$ currencies, where the last currency serves as the numeraire. Customers buy or sell any pair of the $N + 1$

[10] Fuchs, Fukuda, and Neuhann (2023) illustrate some of the limitations. Chaudhary, Fu, and Li (2023) and Davis, Kargar, and Li (2023) estimate the inverse elasticity with respect to securities and then provide a risk-based explanation for their estimates.

currencies. These trades could be motivated by trading demand shocks (e.g., preference shocks) or private information. All customer trades are accommodated by a mass μ of competitive intermediaries. For $n = 1, \dots, N$, the return of currency n from time 0 to time 1 is r_n, which is defined as the return from borrowing one unit of the numeraire at its risk-free rate, converting it to currency n at time 0, investing at currency n's risk-free rate from time 0 to 1, and then converting it back to the numeraire at time 1. We stack r_n into an $N \times 1$ vector as $\mathbf{r} = (r_1, r_2, \dots, r_N)^\top$.[11] Similarly, R_n denotes the return of currency n between time 1 and time 2, which we stack into an $N \times 1$ vector $\mathbf{R} = (R_1, R_2, \dots, R_N)^\top$. We assume there are no redundant currencies, so the matrix $\text{var}(\mathbf{r})$ has full rank, and we assume that the return covariance structure remains stable over time, such that $\text{var}(\mathbf{r}) = \text{var}(\mathbf{R})$.[12] Our goal is to study the price impact of customer trading demand shocks between time 0 to time 1, holding the trading between time 1 and time 2 constant. We empirically map the interval between time 0 to time 1 to a week. Time $t = 2$ represents the long term, where currency prices are no longer affected by trading demand shocks between time 0 and time 1; reaching this stage may take months in reality.

As in the classic asset pricing theory, intermediaries in our model require compensation for any non-diversifiable risks, including those resulting from accommodating customers' trading demand shocks. Analyzing price impacts from trading demand shocks thus provides a way to quantify intermediaries' risk-bearing capacity. The key question is: what risks arise from accommodating trades? If the intermediary is a specialist who only trades one pair of currencies and is perfectly segmented from the rest of the market, then the risks from trading depend solely on the quantity of that pair traded, and the price impact is restricted to that

[11] Throughout this paper, bold font is used to denote matrices and vectors, and $\mathbf{A}^\top$ represents the transpose of $\mathbf{A}$.

[12] All our theory holds if we instead assume the more general form $\text{var}(\mathbf{r}) = L\text{var}(\mathbf{R})$, for some positive constant L.

pair — this resembles the setting of many market microstructure models with a single asset. However, if intermediaries can diversify risks not only within a single currency but also across different currencies — by offsetting trades either internally or through interdealer markets — then the amount of non-diversifiable risks is likely very different from the quantity of any specific currency pair. Consider two customers trading with an intermediary (e.g., a dealer), where one customer is buying and the other is selling. In the simplest case, the two customers are buying and selling the same currency, so the intermediary can offset the risks completely. Even if the customers are trading different currencies, as long as the returns of these trades are not perfectly correlated, the intermediary can still offset some risks. Just as in the classic theory, the ability to diversify across currencies requires the intermediaries to consider portfolios of currencies (i.e., risk factors). Once the trading demand shocks from different currencies are aggregated, the remaining risks are non-diversifiable and the intermediaries require price compensation for bearing such risks.

This portfolio view of trading-induced risk is likely appropriate to describe the FX market, where the intermediaries are highly sophisticated and actively trade among themselves to share risks.[13] Accordingly, we quantify FX intermediaries' risk-bearing capacity with respect to factors that exhibit the most variation in non-diversifiable trading risks (referred to as "traded risk factors"). Section 2.2 shows how to construct these factors. Then, in Section 2.3, we compute the price equilibrium at the factor level to determine the intermediary elasticity of these non-diversifiable trading risks. Finally, when factor prices change, individual currency prices must also adjust to maintain the law of one price in the cross-section; Section 2.4 completes the analysis by computing cross-elasticity, whereby trading demand shocks to one currency generate factor-level non-diversifiable risks, thus affecting factor prices and, in

[13]According to the BIS Triennial Central Bank Survey, about 75% of global FX trades are amongst intermediaries. See also Section 5.

turn, the price of other currencies.

We note that the traded risk factors in Section 2.2 are constructed using observed, equilibrium trading flows and returns. As such, we identify the factors with the largest amount of trading-induced risks in equilibrium. These risks are priced partly due to changes in fundamentals (e.g., the arrival of news, learning from trades) and partly due to intermediaries being pushed against their risk-bearing capacity. Only price responses unrelated to changes in fundamentals can reveal intermediaries' risk-bearing capacity. Hence, once we identify the traded risk factors, we derive the intermediary elasticity of these factors based on hypothetical and marginal trading demand shocks in Section 2.3. Empirically, as trading demand shocks are not directly observed, we instrument for these shocks in our estimation of intermediary elasticity (Section 5).

2.2 Factor Construction

We want to study trading-induced risks that intermediaries bear at the margin. We thus aim to identify a few factors that maximally explain the non-diversifiable risks induced by the aggregate trading flow. Using the U.S. dollar (USD) as the numeraire currency, we first decompose all trades between time 0 and 1 into trades against USD, and express the aggregate trading flow as $\mathbf{f} = (f_1, f_2, \ldots, f_N)^\top$, where f_n is the net customer buying flow for currency n against USD.[14] For any given factor $\mathbf{b}_1 = (b_{1,1}, \ldots, b_{N,1})$, where $b_{n,1}$ represents the weight of currency n in this factor,[15] currency n loads on the factor with a beta $\beta_{n,1} = \mathrm{cov}(r_n, \mathbf{b}_1^\top \mathbf{r})/\mathrm{var}(\mathbf{b}_1^\top \mathbf{r})$. When the intermediary accommodates a currency-level trading flow, f_n, the intermediary effectively bears a factor-level trading flow of size $f_n \beta_{n,1}$,

[14]Specifically, if a customer buys currency n by selling currency m, we record it as a positive trading flow for currency n from USD and a negative trading flow for currency m from USD. In Appendix A.2, we prove that the construction of traded risk factors remains invariant to the choice of the numeraire currency.

[15]By definition, the weight of USD in this factor is $-\sum_{n=1}^{N} b_{n,1}$.

along with other risks uncorrelated with the factor. Given that there are N currencies, intermediaries can offset the factor-level trading flow across different currencies, leaving a non-diversifiable factor-level flow of amount[16]

$$q_1 = \sum_{n=1}^{N} f_n \beta_{n,1} = \text{cov}(\mathbf{f}^\top \mathbf{r}, \mathbf{b}_1^\top \mathbf{r}) / \text{var}(\mathbf{b}_1^\top \mathbf{r}). \tag{1}$$

Note that for any given factor (as defined by the portfolio weights $\mathbf{b}_1$), the factor-level trading flow q_1 varies in proportion to the currency-level trading flow f_n, and their relationship depends on the factor being considered, as varying $\mathbf{b}_1$ changes the beta $\beta_{n,1}$ of a currency to a factor.

We next specify the problem that pins down the most traded risk factors. We note that, because the currency-level flow f_n is measured in the unit of dollars, and beta is unitless, the factor flow q_1 is also measured in the unit of dollars. Multiplying q_1 by the factor return variance $\text{var}(\mathbf{b}_1^\top \mathbf{r})$ therefore changes the unit to trading-induced risks. As our goal is to maximally explain the trading-induced risks using a few factors, we construct the first factor $\mathbf{b}_1$ to maximize the variation of trading-induced risks,

$$\max_{\mathbf{b}_1} \frac{\text{var}(q_1 \text{var}(\mathbf{b}_1^\top \mathbf{r}))}{\text{var}(\mathbf{b}_1^\top \mathbf{r})} = \text{var}(q_1) \text{var}(\mathbf{b}_1^\top \mathbf{r}). \tag{2}$$

We normalize the variance of $q_1 \text{var}(\mathbf{b}_1^\top \mathbf{r})$ by the factor return variance so that scaling $\mathbf{b}_1$ does not affect the objective function.

We construct the second factor $\mathbf{b}_2$ by requiring that the second factor has an uncorrelated return with the first and that the second factor maximizes the variation of trading-induced

[16]Our model assumes a representative intermediary who accommodates all customer trades. In practice, such netting across currencies could also occur through interdealer trading.

risks,

$$\max_{\mathbf{b}_2} \text{var}(q_2)\text{var}(\mathbf{b}_2^\top \mathbf{r}) \qquad (3)$$

$$\text{s.t.}\ \text{cov}(\mathbf{b}_1^\top \mathbf{r}, \mathbf{b}_2^\top \mathbf{r}) = 0,$$

where $q_2 = \text{cov}(\mathbf{f}^\top \mathbf{r}, \mathbf{b}_2^\top \mathbf{r})/\text{var}(\mathbf{b}_2^\top \mathbf{r})$.[17]

Such a sequential maximization procedure bears resemblance to the standard PCA.[18] Because we seek to maximally explain trading-induced risks using a few factors, and these risks depend on both currency-level trading flows and currency returns, our construction is effectively a modified PCA on both trading and returns data. Appendix A.1 provides details on solving for these factors through eigenvalue decomposition. In theory, one can construct at most K factors, where K is the rank of the matrix $\text{var}(\mathbf{f})$. Empirically, a small number of factors are typically sufficient to explain the majority of trading-induced risks.

2.3 Intermediary Elasticity

Having identified the traded risk factors, we now derive the intermediary elasticity of each factor. Our model of the representative intermediary's portfolio optimization is kept deliberately simple to emphasize the relationship between trading and asset prices. We assume that the mass μ of intermediaries have CARA preference.[19] In addition to risk aversion, the only

[17] Because the returns of different factors are uncorrelated by construction, the univariate beta defined here is equivalent to the multivariate beta.

[18] The standard PCA procedure identifies factors for modeling the unconditional expected returns in the cross-section by selecting those that maximally explain the variance of the (unconditional) returns. Specifically, the first factor $\mathbf{b}_1$ maximizes the variance of the factor return: $\text{var}(\mathbf{b}_1^\top \mathbf{r})$. The second factor $\mathbf{b}_2$, conditional on being uncorrelated with the first factor, i.e., $\text{cov}(\mathbf{b}_1^\top \mathbf{r}, \mathbf{b}_2^\top \mathbf{r}) = 0$, again aims to maximize the variance of the factor return: $\text{var}(\mathbf{b}_2^\top \mathbf{r})$, and so on.

[19] We can re-cast the absolute risk aversion as a function of intermediary wealth to mimic a CRRA preference. We empirically investigate the effect of intermediary wealth on intermediary elasticity in Section 5.2.

type of friction that the model features is possible factor-specific frictions in accommodating risks, leading to possibly different factor-specific risk-aversion, denoted by γ_k for factor k.[20]

While the observed factor flows $q_1, \ldots, q_K$ may influence and determine the equilibrium returns in a complex manner, our goal is to compute the sensitivity of factor price change to pure *trading demand shocks*. To achieve this, we examine the price impacts of hypothetical and marginal trading demand shocks $\hat{q}_1, \ldots, \hat{q}_K$, which occur between times 0 and 1 and are uninformed about currency prices at time 2. Due to intermediaries' limited risk-bearing capacity, the price of factor k at time 1 is changed by an extra $\Delta p_k := (P_k(\hat{q}_1, \ldots, \hat{q}_K) - P_k(0, \ldots, 0))/P_k(0, \ldots, 0)$, where $P_k(\hat{q}_1, \ldots, \hat{q}_K)$ is the price of factor k at time 1 after intermediaries accommodate the trading demand shocks.

The equilibrium price impacts are set such that each intermediary finds it optimal to buy $y_k = -\hat{q}_k/\mu$ dollars of factor k. Specifically, as y_k is measured with respect to time-1 currency prices in the absence of these shocks, the time-2 payoff risk from buying y_k dollars of factor k is simply $y_k \mathbf{b}_k^\top \mathbf{R}$. At time 1, these factors are bought at an elevated price $y_k(1 + \Delta p_k)$, which can be compounded to time 2 by multiplying the USD gross risk-free rate R_F. Hence, the representative intermediary's optimization problem reads

$$\{-\hat{q}_1/\mu, \ldots, -\hat{q}_K/\mu\} = \arg \max_{\{y_1, \ldots, y_K\}} \mathbb{E}\left[-\exp\left(-\sum_{k=1}^{K} \gamma_k (y_k \mathbf{b}_k^\top \mathbf{R} - y_k R_F (1 + \Delta p_k)\right)\right]. \quad (4)$$

Applying the first-order condition to (4) and using the assumption that $\mathrm{var}(\mathbf{r}) = \mathrm{var}(\mathbf{R})$, Proposition 1 determines the equilibrium price impact for each factor.

PROPOSITION 1 (Intermediary elasticity). *Denoting $\lambda_k = \gamma_k/(\mu R_F)$, the price im-*

[20] In practice, not all intermediaries may be willing to accommodate risks in every factor. If some intermediaries choose not to absorb risks of a certain factor k, this would manifest as a higher effective risk aversion, γ_k, in our model.

pact of factor k is

$$\Delta p_k = \lambda_k \hat{q}_k \mathrm{var}(\mathbf{b}_k^\top \mathbf{r}). \tag{5}$$

The parameter λ_k is termed the "intermediary elasticity" of factor k. By equation (5), we can express λ_k as follows:

$$\lambda_k = \frac{\Delta p_k}{\hat{q}_k \mathrm{var}(\mathbf{b}_k^\top \mathbf{r})}. \tag{6}$$

Here, Δp_k represents the price impact of factor k at time 1. The denominator, $\hat{q}_k \mathrm{var}(\mathbf{b}_k^\top \mathbf{r})$, measures the change in the quantity of risk due to the marginal trading demand shock into the factor. Consequently, λ_k captures the price compensation that intermediaries require for absorbing the marginal increase in *traded* risk. This concept extends the canonical price of risk that measures price compensation required for taking on an extra unit of *unconditional* risk. Note that in our simple model, λ_k is not a function of intermediaries' pre-existing holdings at time 1, as we do not model nonlinear constraints (e.g., position limits).

We highlight three features of the intermediary elasticity. First, because the traded risk factors have uncorrelated returns by construction, the equilibrium solution from (4) implies that trading demand shocks $\hat{q}_k$ affect only the price of factor k, without influencing any other factors. Appendix A.3 provides a proof. Second, λ_k is invariant to scaling or sign reversal of a factor. This highlights that, economically, $\lambda_k \approx \gamma_k/\mu$ reflects the per-capita risk aversion of intermediaries with respect to that factor.[21] While λ_k is linked solely to γ_k and μ in our stylized model, the empirical estimate of λ_k may also reflect other constraints that intermediaries face when accommodating trading-induced risks for factor k. Third, intermediary elasticity is defined with respect to risks instead of securities. Although quantities of securities are readily observable, in markets where marginal agents optimize their portfolios

[21] Empirically, we can identify λ_k but not γ_k versus μ separately.

to diversify risks, the quantities of risks are more relevant.

2.4 Cross-Elasticity

We now appeal to the law of one price and use factor-specific intermediary elasticity to determine the (inverse semi) cross-elasticity between individual currencies. Consider the scenario where currency m experiences a \$1 trading demand shock (denoted as $\hat{f}_m$ to distinguish it from the overall flow f_m), while customers' trading demand shocks for all other currencies remain constant. As explained in Section 2.2, this additional \$1 trading demand shock for currency m would increase the trading demand shock to factor k by an amount $\beta_{m,k}$, which cannot be further diversified. To induce the intermediaries to bear these additional non-diversifiable factor risks, the price of these traded risk factors must change, as discussed in Section 2.3, which in turn affects the prices of all currencies that load on these risk factors.

Denoting the price impact of individual currency n as[22]

$$\Delta p_n := \frac{P_n(\hat{f}_1, \dots, \hat{f}_N) - P_n(0, \dots, 0)}{P_n(0, \dots, 0)}, \tag{7}$$

where $P_n(\hat{f}_1, \dots, \hat{f}_N)$ is the price of currency n at time 1 after intermediaries accommodate the trading demand shocks $\hat{f}_1, \dots, \hat{f}_N$. Proposition 2 computes the model-implied cross-elasticity.

PROPOSITION 2 (Cross-elasticity). *The cross-elasticity between currencies n and m is:*

$$\frac{\partial \Delta p_n}{\partial \hat{f}_m} = \sum_{k=1}^{K} \frac{\partial \hat{q}_k}{\partial \hat{f}_m} \times \frac{\partial \Delta p_k}{\partial \hat{q}_k} \times \frac{\partial \Delta p_n}{\partial \Delta p_k} = \sum_{k=1}^{K} \beta_{m,k} \times \lambda_k \mathrm{var}(\mathbf{b}_k^\top \mathbf{r}) \times \beta_{n,k}. \tag{8}$$

[22]With a slight abuse of notation, we use Δp_n to denote the price impacts of individual currencies and Δp_k to denote the price impacts of factors.

In the model, only non-diversifiable risks matter; hence, the traded risk factors determine the cross-elasticity, $\partial \Delta p_n / \partial \hat{f}_m$, which is the price impact on currency-n, Δp_n, from currency-m's trading demand shock, $\hat{f}_m$, while holding customers' trading demand shocks into all other currencies constant. Proposition 2 shows that such cross-currency price impacts are channeled via three steps. First, the trading demand shock into currency m changes factor-k trading demand shock $\hat{q}_k$, with the sensitivity given by the beta coefficient $\beta_{m,k}$, as shown in equation (1). Second, changes in factor-k trading demand shock impact its price Δp_k, where the price sensitivity is $\lambda_k \text{var}(\mathbf{b}_k^\top \mathbf{r})$ (Proposition 1). Finally, changes in factor-k price Δp_k impact currency-n price Δp_n through the law of one price, with the sensitivity being $\beta_{n,k}$. Note that Proposition 2 also covers the case where m and n are the same, thereby calculating a currency's price elasticity to its own trading demand shock. Appendix A.4 provides a proof.

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

$$\frac{\partial \Delta p_n}{\partial \hat{f}_n} = \sum_{k=1}^{K} \beta_{n,k}^2 \times \lambda_k \text{var}(\mathbf{b}_k^\top \mathbf{r}) \tag{9}$$

is always positive as long as λ_k is positive. Positive λ_k indicates that intermediaries are averse to bearing trading-induced risks rather than risk-seeking. On the other hand, the cross-elasticity between two currencies could be negative, if the currencies have opposite signs of beta loading to a factor, which reflects complementarity. We return to this point empirically in Section 5.3.

Second, the model-implied cross-elasticity is symmetric between any two currencies n and m, as shown by

$$\frac{\partial \Delta p_n}{\partial \hat{f}_m} = \frac{\partial \Delta p_m}{\partial \hat{f}_n}. \tag{10}$$

This symmetry arises because

$$\frac{\partial \hat{q}_k}{\partial \hat{f}_n} = \beta_{n,k} = \frac{\partial \Delta p_n}{\partial \Delta p_k}. \tag{11}$$

The first equality, relating to the sensitivity of trading demand shock, follows from our portfolio theory (1), while the second equality, concerning price sensitivity, results from the law of one price. Both sensitivities equal the beta of currency n to factor k, which gives rise to the symmetry of the cross-elasticity.

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). CLS provides settlement services to FX trades done by its 72 settlement members, who are mostly large multinational banks.[23] As such, CLS is the largest single source of FX executed data available to the market, covering over 50% of global FX volumes.

We obtain FX order flow data from CLS. Specifically, we have the daily aggregate value of all buy orders and all sell orders done between Banks and their customers in 17 currencies between September 2012 and December 2023. The currencies in our sample include the 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

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

(NOK), New Zealand dollar (NZD), Swedish kroner (SEK), Singaporean dollar (SGD), and South African rand (ZAR). All of our data have Banks as one of the two counterparties in the trade. Trades by Banks encompass trades by dealers who are affiliated with banks and, by extension, trades by hedge funds who trade through their prime brokers. We interpret the trading by Banks as capturing the activities of the representative financial intermediary in our model. The customers in our data, who are Banks' counterparty, are from one of three groups: Funds, which include mutual funds, pension funds, and sovereign wealth funds; Non-bank Financials, which include insurance companies and clearing houses; and Corporate.

To capture the *total* amount of FX risk exposure facing intermediaries, we are the first to jointly analyze the CLS flows data on FX spot with data on FX forward and FX swap. The CLS flows data on FX spot have recently been used in papers that examine topics ranging from market microstructure to the impact of Fed policies (e.g. Ranaldo and Somogyi, 2021; Roussanov and Wang, 2023). Yet as we detail in Appendix Section B, the pronounced negative correlation between flows into spot versus forward and swap means that elasticity estimated from spot flows alone could underestimate the price impact of trading demand shocks. The CLS data on forward and swaps are organized by maturity buckets. We calculate the FX spot exposure inherent in future-settled forward and swaps by discounting the notional using forward rates.[24] From the FX flows data on spot, forward, and swaps, we construct the USD-valued total daily net customer inflow into each currency. As discussed in Section 2.2, we measure all flows relative to USD.

We analyze trading and return at the weekly frequency. We therefore add up daily flows

[24]Specifically, we use the 1-week forward rate to discount back forward and swap contracts with maturity of 1-7 days, the 1-month forward rate for contracts with maturity of 8-35 days, the 3-month forward rate for contracts with maturity of 36-95 days, and the 1-year forward rate for contracts with maturity of greater than 96 days. The choice of forward rate reflects the range of the maturity bucket and forward contract liquidity.

in a week to obtain weekly flows that start every Thursday to the following Wednesday, inclusive. Our final trading data is a panel, between 2012-09-06 and 2023-12-31, 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 data needed to construct currency returns from Bloomberg. To calculate FX returns, we get forward and spot price data for the 16 non-USD currencies in our sample. All prices are at London closing, consistent with our trading flow measure.[25]

We define the weekly currency return as the outcome from borrowing USD at the US risk-free rate today, converting to foreign currency at the spot exchange rate and earning the foreign risk-free rate, then in a week, converting the foreign proceeds back to USD at the future spot rate. That is, for currency n from week t to $t+1$: $r_{t+1,n} = s_{t+1,n} - s_{t,n} + i_{t,n} - i_{t,USD} - x_{t,n} = s_{t+1,n} - f_{t,n}$, where s is the log spot rate, f is the log forward rate, i is the net risk-free rate, and x measures the deviation from the covered interest-rate parity (CIP). Throughout, we define exchange rates as the number of USD per one unit of foreign currency; a higher s thus corresponds to a depreciation of USD. Note that our currency return includes CIP deviation ($x_{t,n} = f_{t,n} - s_{t,n} - i_{t,USD} + i_{t,n}$), so as to capture compensation for all risks that intermediaries take in absorbing customer flows, including possible inventory costs arising from balance sheet constraints.

[25] CLS records daily flow as all orders submitted during the FX business day, which follows the London FX market hours.

3.3 Other Data

We collect data on various sovereign bond auctions to instrument for FX trading demand shocks. Specifically, we obtain from government websites the auction announcement data for the U.S. Treasury bond auctions, the Australian Treasury bond auctions, the Canadian Treasury bond auctions, the U.K. Gilt auctions, the Japanese government bond auctions, the Italian government bond auctions, the French OAT auctions, and the German Bund auctions.

We also collect various data to construct excess returns in six other asset classes. 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. Returns to these CDS indices are defined from the perspective of the seller. For commodities, we obtain six Bloomberg commodity futures return indexes on energy, grains, industrial metals, livestock, precious metal, and softs. For corporate bonds, we obtain five Bloomberg indices on U.S. corporate bonds by credit rating (Aa, A, Baa, high yield; we exclude AAA to avoid collinearity with the risk-free rate). For equities, we use the "Market" return from Ken French's website, which is the value-weighted returns from all publicly traded U.S. firms in CRSP. For options, we obtain call and put pricing data on S&P500 from OptionMetrics, and construct portfolios of leverage adjusted option returns following Constantinides, Jackwerth, and Savov (2013). For US Treasury bonds, we get yields of the six maturity-sorted "Fama Bond Portfolios" from CRSP. We exclude the portfolio of 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 to calculate excess returns. The Bloomberg CDS data are available from 2007 onward, and the OptionMetrics data are available until December 2022. Data on all other asset classes start in January 2000 and end in December 2023.

4 Traded Risk Factors in FX

In this section, we identify the most important 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 show that these risk factors can be interpreted as the Dollar, the Carry, and the Euro-Yen, respectively. Finally, we highlight that these risk factors also capture the preponderance of the unconditional return variations in individual currencies.

4.1 Baseline Traded FX Factors

Our objective is to find risk factors that can capture FX trading's impacts on currency prices in the cross-section. We therefore look for factors that maximally explain trading-induced risks. Following the procedure outlined in Section 2.2, we identify the traded risk factors using weekly net flows ($\mathbf{f}$) and log returns ($\mathbf{r}$) of 16 non-USD currencies. 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.[26] For example, in Factor 1, for every $1 bought, $0.15-worth of CAD and $0.5-worth of EUR are sold. As discussed in Section 2.3, we can freely scale each factor without affecting the intermediary elasticity, and we accordingly scale all factors to facilitate comparison.[27] Because the risk factors that we identify are traded, they logically place more weight on currencies that are more widely traded. In particular, six developed economy currencies have consistently high weights across the top 3 factors, these currencies are AUD, CAD, CHF, EUR, GBP, and JPY, and we highlight them in red. Of the total trading-induced

[26]The portfolio weight of USD is the negative sum of the weights of all other currencies.

[27]Specifically, 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.

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 outlined in Section 2.2. We use weekly return and flow data for 16 non-USD currencies from September 2012 to December 2023. The portfolio weight of USD is computed as the negative sum of the weights of all other currencies.

non-diversifiable risks, given by $\sum_{l=1}^{K} \text{var}(q_l)\text{var}(\mathbf{b}_l^\top \mathbf{r})$, our three traded risk factors account for 65%, 16%, and 9%, respectively. In other words, these three factors explain approximately 90% of the risk that intermediaries bear when accommodating trading flows.

Principal component analysis (PCA) is often seen as sensitive to minor changes in data. Yet the traded FX factors identified through our modified PCA procedure are robust to changes in the sample period. In Table 2, we report the correlation between the traded FX factors identified using our modified PCA on the full sample and those identified using

Table 2: **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.

the pre- and post-2020 subsamples. The correlations for both returns and flows are notably high, approaching 1 for the first factor and exceeding 0.8 for the other two factors. This evidence suggests that the underlying data are well-behaved, and in particular, that the flow and return covariance structures are rather stable over time.

A tempting alternative approach may be to find traded FX factors by performing a PCA directly on trading data. Portfolios from such an approach would simply place weight on one single major currency. For example, as Appendix Table A3 illustrates, the first such "traded FX factor" would place a portfolio weight of -1 on EUR and 0 on all other non-USD currencies, reflecting that EUR/USD is the most actively traded pair. This result arises because the standard PCA simply points to the portfolios with the largest trading in dollar amounts, whereas the portfolios in Table 1 identify those with the greatest trading-induced risks.

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, respectively. Examining Factor 1 in Table 1, we see that all non-USD currencies enter the portfolio with a negative weight. This pattern is reminiscent of the proverbial Dollar portfolio, which shorts all non-USD currencies simultaneously and bets on the USD exchange rate. We thus propose a traded Dollar factor that goes long in USD and shorts the six most traded currencies (AUD, CAD, CHF, EUR, GBP, and JPY) in equal weights. In contrast, Factor 2 in Table 1 has positive weights on high interest rate currencies, e.g., AUD, and negative weights on low interest rate currencies, e.g., JPY. This pattern coheres with the proverbial Carry portfolio, which bets on violations of the uncovered interest-rate parity (UIP). We accordingly propose a traded Carry factor that goes long in AUD, CAD, and GBP in equal weights, and shorts CHF, EUR, and JPY in equal weights. Finally, Factor 3 in Table 1 has a large positive weight on EUR and a large negative weight on JPY. We postulate a traded Euro-Yen factor that goes long in EUR and shorts JPY. This Euro-Yen factor reflects the active currency trading between two of the world's largest economies, the Euro area and Japan, even as the Dollar and Carry factors place EUR and JPY on the same side of trading. In other words, a long-short position in EUR-JPY generates non-diversifiable risks, even after hedging out the Dollar and Carry factors.

The data support our interpretation of the traded FX factors. Using our proposed Dollar, Carry, and Euro-Yen factor weights, we construct factor returns and factor flows.[28] In Table 3, we show the correlation between the baseline traded FX factors in Table 1 ("PC

[28]Specifically, we perform the procedure outlined in Section 2.2 after projecting the currency-level returns and flows onto the space spanned by the proposed Dollar, Carry, and Euro-Yen factor weights. This procedure ensures that resulting Dollar, Carry, and Euro-Yen factors have uncorrelated returns with each other.

Table 3: **Correlation between Return and Flow to Baseline PC Factors versus to 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 displays the correlation between return and flow to baseline traded FX factors in Table 1 ("PC Factors") and return and flow to traded FX factors constructed from the proposed factor weights of the Dollar, the Carry, and the Euro-Yen ("Economic Factors").

Factors") and the traded FX factors constructed from the proposed factor weights ("Economic Factors"). The correlations are close to 1 in both returns and flows for all three factors. Together, the three Economic Factors can explain about 86% of all trading-induced non-diversifiable risks, close to the risks accounted for by the PC Factors. Given the striking similarity between the PC Factors and the Economic Factors, we focus on analyzing the more interpretable Economic Factors in the rest of the paper.

The construction of traded FX risk factors reveals two important results. First, that just three factors can explain almost all trading-induced risks validates our conceptual framework. Indeed, we postulate that non-diversifiable trading risks are what get priced by intermediaries; in such a world, risks induced by FX trading would exhibit a strong factor structure. Second, the traded FX factors have clear economic interpretations, elevating the relevance of our analysis on intermediaries' risk exposure and risk-bearing capacity. For example, the Carry trade is a popular FX trading strategy and the Carry factor is a well-known risk factor in FX returns (e.g., Lustig, Roussanov, and Verdelhan, 2011). Yet it is not previously known how much Carry trades are done[29] or how trading shocks to the Carry affect its returns. We answer the first question in this section by estimating intermediaries' cumulative exposure to the traded risk factors. We tackle the second question in Section 5.2 by estimating the intermediary elasticity of different traded FX factors.

The economic exposure to, say, the Carry trade is challenging to assess from trading data alone: FX traders can simultaneously buy and sell multiple currencies, and exposures in one currency can be quickly hedged or diversified away by trading other currencies. Yet the traded FX risk factors represent precisely what is non-diversifiable after accounting for trading in all currency pairs. Cumulative flows into the risk factors thus show the cumulative exposure to the currency portfolios that define the risk factors.

Figure 1 plots the cumulative flow to each of the three factors by customer type. As detailed in Section 3.1, there are three types of customers: Funds; Corporates; and Non-Bank Financials. We also plot in dashes the Net Total, which represents the cumulative net customer flows that Banks need to absorb. By market clearing, the *negative* of the Net Total represents the intermediaries' cumulative flow. Panel (a) illustrates the flow to the Dollar factor. Over the sample period, Funds are persistently selling Dollars, whereas Corporates are persistently buying. In recent years, the buying pressure has been so strong that intermediaries have been net sellers of the Dollar factor. We note that intermediaries, especially dealers, may not be able to maintain a sustained inventory imbalance. The provision of USD here likely reflects positions of hedge funds and/or USD deposits or wholesale funding made available by (dealer affiliated) banks (Du and Huber, 2024). Panel (b) illustrates the flow to the Carry factor. We observe that customers had not taken large directional bets with the Carry factor in the first half of our sample, but since then — especially post-2022 —

[29] https://www.economist.com/leaders/2024/08/15/time-to-shine-a-light-on-the-shadowy-carry-trade

Figure 1: **Cumulative Flow by Investor Type to Top 3 Traded FX Factors**

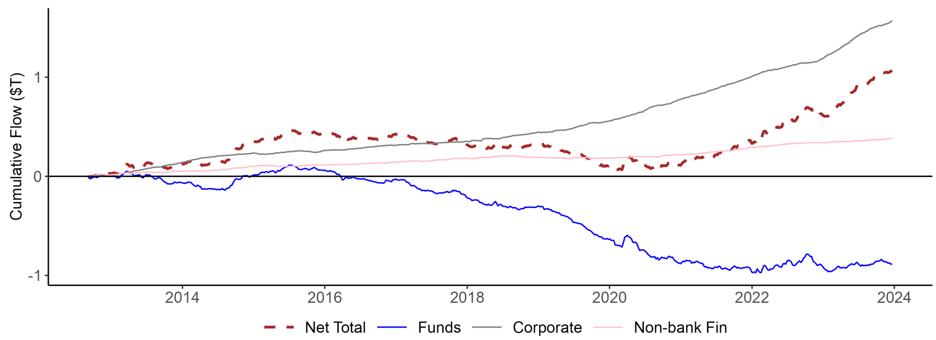

(a) **Dollar Factor**

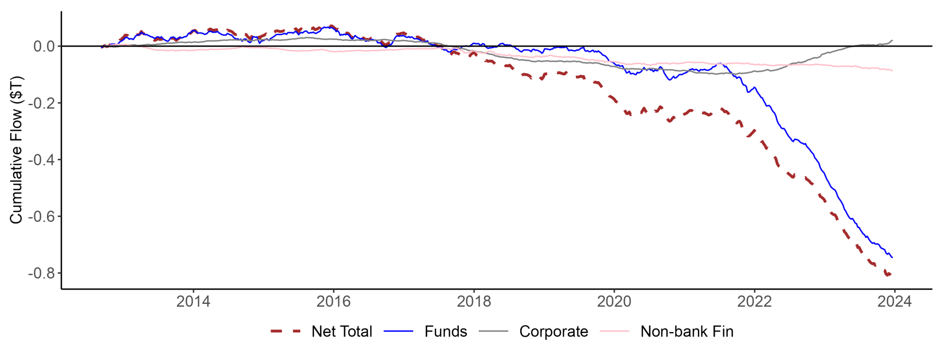

(b) **Carry Factor**

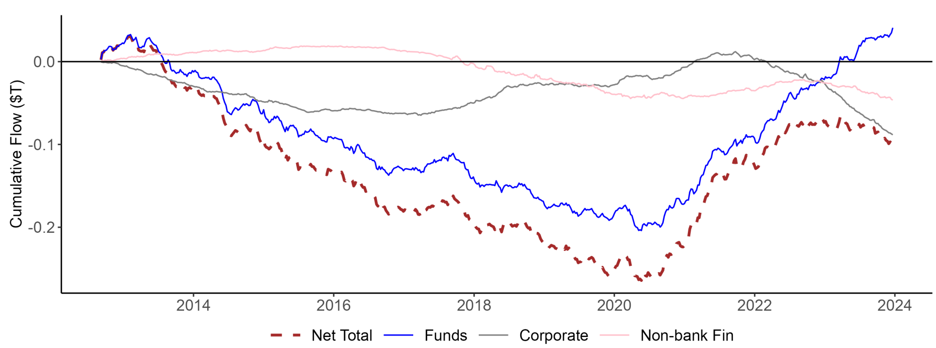

(c) **Euro-Yen Factor**

Notes: This figure displays the cumulative flows of the top three traded FX factors over our sample period, from September 2012 to December 2023. There are three types of customers: Funds, Corporates, and Non-Bank Financials. In addition to these customer flows, we plot in dashes the Net Total, which represents the net customer flows that Banks (intermediaries) need to absorb.

customers (in particular Funds) have sold off the Carry factor. As a result, the Carry trade exposure borne by Banks, or intermediaries including dealers and hedge funds, accumulated to $0.8 trillion between 2012 and 2023. Finally, from Panel (c), we see that the Euro-Yen factor was sold by both Corporate and Funds right up to around the Covid-19 Crisis in 2020. Thereafter, Funds have bought back all of their short positions to approximately neutral, while Corporate continued to sell the Euro-Yen factor. Thus, intermediaries have also been accumulating exposures in the Euro-Yen. As JPY serves as a "funding currency" (negative weight in the portfolio) in both the Carry and the Euro-Yen, our analysis underscores that an unwinding of intermediaries' short JPY positions may not simply be a story about Carry trade.

In Figure 2, we plot the cumulative returns to the three traded FX factors. The Dollar return has been strong, as USD has appreciated against most currencies over the last decade. The Carry return had been around zero until 2022, but has since taken off. The uptick in the Carry return coincided with the rapid accumulation of Carry exposures by intermediaries. This suggests that FX trading may affect the pricing of these traded risk factors. We return to this in Section 5.

4.3 Factor Decomposition of Individual Currencies

Although the three traded FX factors are designed to maximally capture *trading-induced* risks, we show that these factors also explain a substantial amount of *unconditional* return of individual currencies. Figure 3 illustrates the decomposition of individual currency's return into the Dollar, the Carry, and the Euro-Yen factors. This decomposition is achieved by regressing currency-level returns on the returns of traded FX factors in the time series

Figure 2: **Cumulative Return of Top 3 Traded FX Factors**

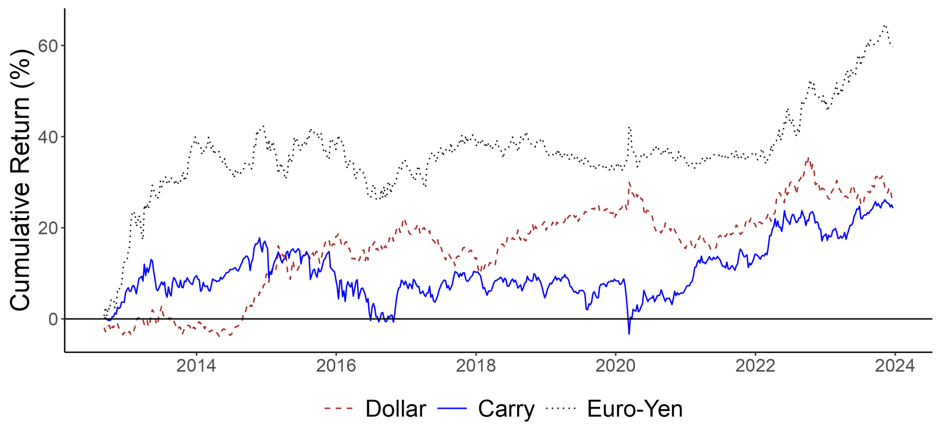

Notes: This figure displays the cumulative returns of the top three traded FX factors over our sample period, from September 2012 to December 2023.

and examining the contribution of each factor to the R^2.[30] The Dollar, the Carry, and the Euro-Yen factors together account for between 69% and 94% of individual currency's unconditional return. The fact that these traded risk factors can explain a high fraction of individual currency's return variation is not a foregone conclusion: the traded FX factors are only designed to rationalize changes in FX risk that are induced by trading, not necessarily all the risks that are priced in the FX markets.

Because individual currencies load meaningfully on the traded FX factors, trading demand shocks to these factors can have significant impacts on currency prices. Moreover, as there are substantial cross-sectional differences in both the direction (illustrated with plus and minus signs) and the magnitude (proportional to the square root of the size of the bar) of individual currencies' factor loading, the cross-substitution patterns among currencies are potentially complex. We map out currency-level cross-elasticity in the next section.

[30]Because the returns of different factors are uncorrelated by construction, the regression R^2 from each factor is additive.

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

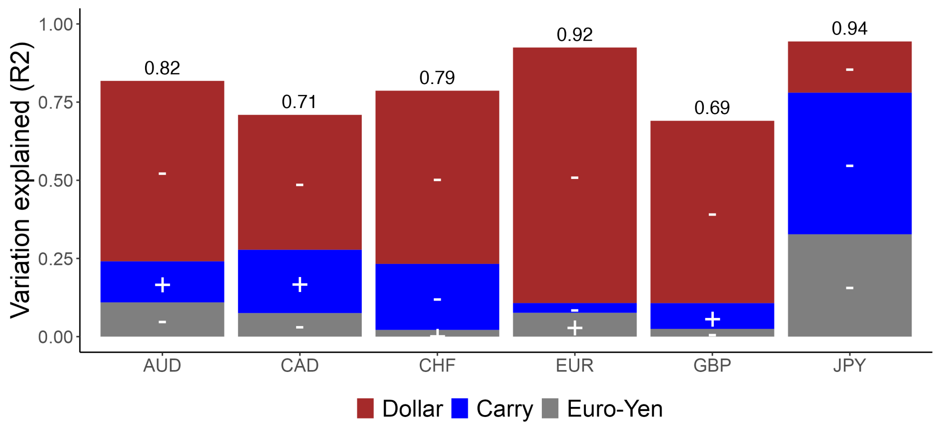

Notes: This figure decomposes the returns of individual currencies into the Dollar, the Carry, and the Euro-Yen factors. The decomposition is achieved by regressing currency-level returns against the returns of the traded FX factors in the time series. It plots the marginal R^2 values attributed to each factor and labels the total R^2. The positive and negative signs illustrate the direction of the beta loadings.

5 Elasticity in FX Markets

In this section, we estimate intermediary elasticity of the traded FX risk factors and apply the estimates to recover the own- and cross-elasticity between an entire panel of individual currencies.

5.1 Instrument Construction

We are interested in estimating λ_k, the intermediary elasticity of traded FX factor k, in equation (5). Because the traded FX factors are constructed to have uncorrelated returns, we apply Proposition 1 to estimate λ_k factor-by-factor without worrying about any cross-factor substitution. For each factor, however, we must instrument for the unobserved trading demand shocks. As discussed in Section 2, the observed trading flows q_k could affect factor prices in a complex manner. Because we want to estimate the sensitivity of price to only trading-induced risks, we must instrument for trading demand shocks $\hat{q}_k$ that are orthogonal to changes in fundamentals.

For each factor $k = \text{Dollar}, \text{Carry}, \text{Euro-Yen}$, we want to run a time-series regression of the factor's weekly observed return $r_{k,t}$ on its instrumented weekly flow $\hat{q}_{k,t}$:

$$r_{k,t}/\text{var}(r_{k,t}) = \lambda_k \hat{q}_{k,t} + \epsilon_{k,t}, \text{ where} \tag{12}$$

$$q_{k,t} = \theta_k z_{k,t} + e_{k,t}, \tag{13}$$

$$\text{cov}(z_{k,t}, \epsilon_{k,t}) = 0. \tag{14}$$

The instruments (z_k) for the observed factor flows (q_k) must be both relevant (equation (13)) and valid (equation (14)). We propose sovereign bond auction announcements as instruments. Government entities such as the US Treasury periodically auction off long-term debt obligations, e.g., US Treasury notes and bonds. Foreign investors participate in advanced economies' auctions. For example, foreign investors directly purchased on average 14% of US Treasury notes and bonds sold at auctions between September 2012 and December 2023.[31] Because foreign investors need to exchange their domestic currencies for the local currency of the auction to participate, the instruments have relevance.

We moreover argue that these auctions, and in particular, the *announcements* of the offered amount at upcoming auctions are valid. First, auction announcements are plausibly exogenous to fundamentals in $\epsilon_{k,t}$ because auctions are often forward-guided. For example, each December, the Finance Agency of Germany releases the auction calendar for the upcoming year, including the target auction amount at each auction. In other words, the news content of auction announcements is likely limited. Indeed, Wachtel and Young (1990) find

[31] Foreign investors' actual purchase of auctioned Treasury securities could be higher, as this 14% excludes foreign purchases done indirectly via U.S. investment funds and dealers.

that while Treasury auction *results* move bond yield, the *announcements* have no detectable effect. Second, auction announcements likely satisfy the exclusion restriction because, unlike auction results, they do not reflect investors' actual purchase and, therefore, do not reveal investors' potentially private information about fundamentals. In other words, intermediaries' price response to the auction-announcement-instrumented trading is likely driven by their risk-bearing capacity and not by any information update.

As the traded FX factors that we need to instrument load on multiple currencies, we consider sovereign auction announcements from a panel of countries. Specifically, we use the US Treasury auction announcements as the instrument for the Dollar factor. We use the Australian, Canadian, British, and Japanese government bond auction announcements as the instruments for the Carry factor. We use the announcements of Euro-Area Government bond auctions, defined as the sum of German, French, and Italian government bond auctions, as the instrument for the Euro-Yen factor. For each auction, we aggregate the offered amount across all announcements in a week, as in FX trading flows.[32] To instrument for factor flows in week t, we use announcements in week t for the Dollar and the Carry, and announcements in weeks $t-1$ and t for the Euro-Yen. The longer window for European sovereign auctions allows for potential delays in auction-induced currency conversion, which is a relevant concern because sovereign auctions in Germany, France, and Italy do not allow direct bids from foreign investors. Finally, we remove any linear trend in the size of the auctions over time.

Table 4: **Estimated Intermediary Elasticity**

	Dollar		Carry		Euro-Yen	
	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)
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 factors, based on regression (12). 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 from 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. *$p<.1$; **$p<.05$; ***$p<.01$.

5.2 Factor-Level Elasticity

Table 4 presents the λ_k estimation results for the Dollar, Carry, and Euro-Yen factors. For all three factors, the estimated intermediary elasticity is positive and statistically significant, pointing to intermediaries having limited risk-bearing capacity. Both the OLS and the IV estimates show that the intermediary elasticity is the smallest for the Dollar and the largest for the Euro-Yen. In other words, intermediaries are best able to bear marginal risks in the Dollar factor, and their risk-bearing capacity declines for the Carry and is even less for the Euro-Yen. We note that the OLS estimates are similar to the IV estimates but slightly smaller. The primary bias mitigated by the instrument is the correlation be-

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

Table 5: **Economic Magnitude of Intermediary Elasticity**

	Int. elasticity (IV) λ_k	Return volatility $\sigma(r_{k,t})$	Price impact per \$B trading $\lambda_k \sigma^2(r_{k,t})$
Dollar	0.11	6.9%	5.0 bps
Carry	0.14	8.2%	9.3 bps
Euro-Yen	0.34	9.4%	29.3 bps

Notes: This table interprets the economic magnitude of intermediary elasticity for the Dollar, Carry, and Euro-Yen factors. The columns report, from left to right, the elasticity estimates from the IV regression, the standard deviations of factor returns, and the impact on factor prices per billion of trading demand shock.

tween information-driven price changes $\epsilon_{k,t}$ and contemporaneous customer flows q_k. This correlation is negative because customers trade against fundamentals: when news causes a currency to depreciate, customers buy, and vice versa. Such behavior is consistent with the profitability of momentum strategies in FX (Menkhoff, Sarno, Schmeling, and Schrimpf, 2012).[33]

Table 5 interprets the economic magnitude of the estimated intermediary elasticity λ_k. Recall that λ_k is defined with respect to a unit of risk. We use Proposition 1 to calculate the price impact of a \$1 billion factor-level trading demand shock on factor price as $\lambda_k \sigma^2(r_{k,t})$, where $\sigma(r_{k,t})$ is the annualized volatility of the factor return. Our estimates imply that a \$1 billion flow into the Dollar, Carry, and Euro-Yen factors increases their price by 5, 9, and 29 basis points, respectively.

These price impacts are large compared to the estimated (inverse) elasticity of US equities. Gabaix and Koijen (2021) estimate that a \$1 billion trading demand shock to the entire U.S. stock market raises the aggregate price by about 1.7 bps.[34] A shock to the en-

[33] In a rational market, prices would adjust to fundamental news without trading (Milgrom and Stokey, 1982). In reality, when customers buy in response to negative fundamental news, they cause prices to under-react, leading to subsequent price drift in the same direction, which generates price momentum.

[34] Gabaix and Koijen (2021) find that a 1% greater trading demand shock in the entire US stock market

tire U.S. market is a shock to the non-diversifiable market factor and is thus comparable to intermediary elasticity in FX. The price impact of a comparable shock is much larger in FX, highlighting the much more limited risk-bearing capacity in FX. Viewed through Proposition 1, the more limited risk-bearing capacity in FX (higher intermediary elasticity) reflects either a greater aversion to FX risks (γ_k) or a smaller amount of arbitrage capital (μ). We think that it is likely that the supply of FX arbitrage capital is limited because the FX market is highly specialized, where only sophisticated participants like bank dealers and hedge funds offer to absorb trading demand shocks. This may seem counterintuitive given the large turnover in FX, but as much as 75% of all FX trades are among intermediaries (BIS, 2022), suggesting that the amount of arbitrage capital available to absorb shocks is much smaller than the total turnover.[35] We also note that the cross-factor variation in the estimated intermediary elasticity could reflect differences in available arbitrage capital across risk factors, whereby lesser-known and less-traded factors such as the Euro-Yen have more limited arbitrage capital.[36]

Our estimates can be interpreted to reflect the average intermediary elasticity outside of crisis periods. We exclude the first half of 2020 from our estimation because markets experienced extreme price volatility and dislocation during this period, casting doubts over our instruments' validity and strength. Yet, even outside of crisis periods, it is possible that intermediaries' risk-bearing capacity varies with their wealth or regulation-induced balance-sheet constraints. In Appendix Table A4, we explore the state-dependency of intermediaries'

increases price by 5%. The average market capitalization between 2012 and 2022 is about \$31.7 trillion. A \$1 billion trading demand shock therefore raises the price of the market factor by 1.7 bps over our sample period.

[35] Of the FX trades accounted for in the BIS Triennial Central Bank Survey, 46% are between reporting dealers, 22% are with non-reporting dealers, 7% are with hedge funds. All these entities are captured in Banks in our data and correspond to the intermediaries in our model.

[36] The annualized volatility of customer flows is \$85 billion for the Dollar factor, \$34 billion for the Carry factor, and \$22 billion for the Euro-Yen factor.

risk-bearing capacity by interacting the instrumented trading demand shocks in the most traded Dollar factor with proxies of wealth or balance-sheet constraints.[37] Specifically, we use the cumulative weekly return of the KBW Bank Index to capture variations in intermediaries' wealth,[38] and we use the 3-month AUD-JPY cross-currency basis to capture variations in regulation-induced balance-sheet constraints.[39] The interaction terms with equity return and with CIP deviation both have intuitive signs: when intermediaries' wealth is higher, their risk-bearing capacity likely increases, and the intermediary elasticity decreases; in contrast, when CIP deviation opens up, intermediaries are more constrained, which likely decreases their risk-bearing capacity and increases intermediary elasticity. However, these two interaction terms are not statistically significant. We conclude that there is suggestive evidence of state-dependent risk-bearing capacity, but our empirical setting does not have the power to establish this variation.

We note that, also in Appendix Table A4, the usage of the Federal Reserve's central bank liquidity swap lines *does* affect the intermediary elasticity for the Dollar factor.[40] A one standard deviation increase in swap line usage statistically significantly decreases the intermediary elasticity for the Dollar factor by 0.056, which is approximately 52% of the

[37]To implement interaction in the context of instrumented flow, we run two first-stage regressions, one for factor flow, and one for factor flow interacted with the time-varying measure. Both regressions include the instrument, the time-varying measure, and the interaction between the instrument and the time-varying measure as explanatory variables. We demean and standardize the time-varying measure.

[38]KBW Nasdaq Bank Index is the value-weighted average of 24 banking stocks representing the largest U.S. national money centers. This series has a 0.86 correlation with the intermediary equity wealth constructed using primary dealers' bank holding company in He, Kelly, and Manela (2017).

[39]Cross-currency basis measures deviations from covered interest-rate parity (CIP) and captures regulatory risks that prevent intermediaries from fully taking advantage of investment opportunities (Du, Hébert, and Huber, 2022). The basis between AUD-JPY is the largest among developed country currency pairs. We use the average basis in week t with factor flow in the same week.

[40]Swap lines were set up during the financial crisis of 2007-09, and have been used throughout our estimation period to provide occasional dollar funding to foreign central banks who then pass the funding to local intermediaries. Given the friction in dispersing funding to local intermediaries, we use swap usage in week $t-1$ with factor flow in week t. The Fed often conducts small-value operations. We exclude them but the results are invariant to their inclusion.

average elasticity (0.107) estimated in column (2) of Table 4. Because the swap line provides immediate dollar funding to intermediaries, we interpret its effect on the Dollar factor's intermediary elasticity to mean that USD funding constraints constitute a key reason why the Dollar factor is priced.

Finally, we note that the precision of an IV estimation depends on the strength of the instrument. The heteroscedasticity and autocorrelation consistent (HAC) effective F-statistics of the instruments are 24.8, 6.5, and 3.8, respectively, for trading flows into the Dollar, the Carry, and the Euro-Yen. The effective F-statistics for the Carry and the Euro-Yen are below the rule-of-the-thumb threshold of 10. To better understand the implications of using potentially weak instruments on the 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, the Anderson-Rubin confidence interval is bounded away from zero, but is very wide in the other direction. In other words, we are reasonably confident that the instrumented intermediary elasticity is not zero; however, we are much less certain that the true value is not larger. A larger estimate would mean even greater price response for a given unit of risk, implying even more limited risk-bearing capacity.

5.3 Cross-Currency Elasticity

Based on the IV estimated intermediary elasticity λ_k, we apply Proposition 2 to compute the cross-currency elasticity and report the results in Table 6. 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. Each entry shows the price response in one row (column) currency, in basis points, to a $1 billion trading demand shock in the corresponding column

Table 6: **Cross-Currency Elasticity**

	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 2, the estimated factor-level elasticity λ_k from Table 4, and the beta loadings of currencies to factors (signs illustrated in Figure 3) to compute cross-currency elasticity. Each entry represents the price movement in bps of a row (column) currency, as induced by a $1 billion trading demand shock into a column (row) currency. As noted after Proposition 2, the model-implied cross-elasticity is symmetric, so we report only the upper half.

(row) currency. For instance, the entry of 7.9 in the first row and second column implies that a \$1 billion trading demand shock to the CAD (AUD) increases the return of AUD (CAD) by 7.9 bps, holding the trading demand shocks in all other currencies equal. As noted after Proposition 2, our model-implied cross-elasticity is symmetric because we are capturing the cross-currency price impact as channeled via the three most traded FX factors.

Table 6 reveals several interesting patterns of cross-currency elasticity. First, all entries are positive. This is because all currencies load on the Dollar factor in the same direction, which is the most important traded risk factor in the cross-section. Second, the cross-elasticity 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. This modest cross-elasticity owes to opposite beta loadings with respect to the Carry factor. In effect, currencies in one of these two groups hedge currencies in the other group in risk exposures to the Carry factor. In IO, such phenomena are typically referred to as complementarity. Third, we note that although EUR and JPY are both low interest-rate currencies, we estimate a rather small cross-elasticity because the two currencies are on the opposite side of the Euro-Yen factor. This result suggests that EUR and JPY are not entirely substitutible.

Moreover, although we analyze traded FX factors constructed based on the six major currencies and USD, we recover meaningful cross-elasticity in other currencies due to these currencies' loadings on the three traded FX factors. As a sanity check of our methodology, we examine the cross-elasticity for HKD, a currency pegged to USD within a narrow band of 1%. We do not use this pegged information in our estimation. We observe that the entire column and row associated with HKD are close to zero. The minimal impact vis-à-vis all other currencies reflects the nature of a pegged currency, whose own trading demand shocks have negligible risk implications for any currencies, and whose exchange rates relative to

USD are not meaningfully impacted by trading demand shocks in any other currencies.

By estimating a panel of currency-level cross-elasticity using factors, we circumvent two key estimation challenges. First, trading demand shocks likely correlate across currencies. Finding cross-elasticity via regressions directly at the currency level would require instrumenting the trading demand shocks in every currency and will likely have low power due to multicollinearity. Second, the substitution patterns among currencies may be complex. Models in industrial organizations typically exploit user demographics or product characteristics to inform such substitution, but such information may not be applicable to financial assets that are valued for their risk-return profiles. Our approach appeals to the law of one price to reduce the many currency-level cross-elasticity to just three factor-level intermediary elasticities, λ_k, which we carefully estimate using instrumental variables. As individual currencies' loading on the factors can vary in magnitude and in direction, when we map the estimated λ_k to the currency-level cross-elasticity, we recover a rich and nuanced cross-substitution pattern.

6 Intermediary Elasticity and Cross-Asset Elasticity

In this section, we use the traded FX risk factors to inform cross-elasticity between asset classes. If intermediaries arbitrage across markets, then trading demand shocks in any one market could impact the prices in all markets by affecting the commonly priced traded FX factors. We show that six non-FX assets indeed load on the traded FX factors, and use the intermediary elasticity of the trade FX factors to compute between-asset cross-elasticity.

Figure 4: **Cross-Asset Return Variation Explained by Traded FX Factors**

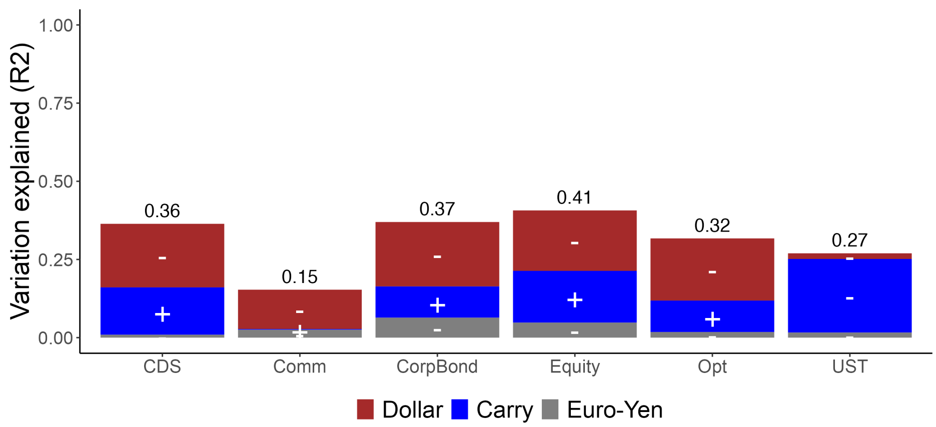

Notes: This figure decomposes the returns of individual assets into the Dollar, the Carry, and the Euro-Yen factors. The decomposition is achieved by regressing asset class monthly average excess return between 2000-02 and 2023-12 on returns from the three traded FX factors. The returns from CDS are available starting in 2007-04. The returns from Opt end in 2022-12. The figure 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.

6.1 Other Assets' Loading on Traded FX Factors

We study six non-FX asset classes: credit default swap (CDS), commodities (Comm), corporate bonds (CorpBond), equities (Equity), equity options (Opt), and US Treasury bonds (UST).[41] We regress each asset class' monthly average excess return between 2000-02 and 2023-12 on returns from the Dollar, the Carry, and the Euro-Yen. We use the R^2 from these regressions to measure the amount of asset return variations explained by the three traded FX factors and illustrate the results in Figure 4.[42] Figure 4 also illustrates the sign of asset class m's return loading on each of the three traded FX factors ($\beta_{m,k}$ in Proposition 2).

The three traded FX factors jointly explain between 15% (commodities) and 41% (eq-

[41]We construct the return of each asset class as the equal-weighted average return of all available portfolios.

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

uities) of the returns in the six non-FX asset classes we examine. In other words, markets are neither fully integrated, where only non-diversifiable risks that are systematic across all markets are priced by intermediaries (extending Sharpe (1964) to the representative intermediary in He and Krishnamurthy (2013)), nor completely segmented such that risks priced in each market are idiosyncratic to that market (possibly due to reasons in Siriwardane, Sundaram, and Wallen (2022)). We note that our analysis provides a lower bound on the degree of integration between markets, as there could be non-traded common risk factors priced across markets.[43]

Figure 4 shows that each asset class has its unique loading on the three traded FX factors, both in terms of magnitude and in terms of direction. To start, while the Dollar factor is statistically significantly present in the return of all six asset classes, it is least important in explaining the return of US Treasury bonds.[44] Moreover, while all other asset classes load positively on the Carry factor, the US Treasury bonds load negatively on it. This contrast suggests that large shocks to the Carry factor could be a reason for divergent price movements in US Treasury bonds versus other assets. Finally, the Euro-Yen factor is less prominent in non-FX asset classes but it does explain a non-negligible amount of return in corporate bonds and equities.

[43]We also explore the explanatory power of traded FX factors for other assets' returns outside of crisis periods (e.g., GFC, Covid). As Appendix Figure A3 shows, the results are largely similar.

[44]Given that foreign investors hold nearly a quarter of Treasury bonds, and that their demand potentially affects both Treasury returns and exchange rate, it may be surprising that Treasury returns load so little on the Dollar factor. 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).

Table 7: **Asset Elasticity to Traded FX Factors**

	CDS	Comm	CorpBond	Equity	Opt	UST
Dollar	-2.0	-5.0	-2.8	-4.4	-4.4	-0.5
Carry	3.7	1.6	3.7	7.9	6.1	-2.3
Euro-Yen	-2.5	-10.3	-7.3	-10.8	-6.6	-1.6

Notes: This table uses Proposition 2, the estimated factor-level elasticity λ_k from Table 4, and the beta loadings of assets to factors (signs illustrated in Figure 4) to compute cross-elasticity 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 trading demand shock into a traded FX factor.

6.2 Cross-Asset Elasticity

The presence of traded FX factors in other asset classes implies that, because of intermediaries' limited risk-bearing capacity, trading in FX could also affect prices in these asset classes. In Table 7, we report the price impact in non-FX assets due to a $1 billion trading demand shock to each of the three traded FX factors. The magnitude of the price impact depends on two forces: the loading of an asset class m on traded FX factor k, and the intermediary elasticity of factor k. Because intermediaries have more risk-bearing capacity for the Dollar factor, the price impact from a $1 billion trading demand shock to the Dollar factor is rather modest, even though most asset classes load heavily on this factor. In contrast, for a shock of the same size, the Carry factor and the Euro-Yen factor elicit much stronger price responses in other asset classes. However, we note that shocks to the Carry and the Euro-Yen factors are likely smaller in magnitude: the annualized volatility of customer flows for the Dollar factor ($85 billion) is much higher than that for the Carry ($34 billion) and the Euro-Yen ($22 billion).

We can moreover consider cross-market elasticity as channeled through the traded FX factors. Trading demand shocks in any one market would alter the intermediary's exposure to

Table 8: **Cross Elasticity Between Assets Due to Traded FX Factors**

	CDS	Comm	CorpBond	Equity	Opt	UST
CDS	2.4	3.5	3.2	5.8	4.7	-0.5
Comm		8.9	6.0	9.5	7.7	0.7
CorpBond			4.8	8.3	6.5	-0.2
Equity				14.6	11.4	-0.9
Opt					9.3	-0.6
UST						0.7

Notes: This table uses Proposition 2, the estimated factor-level elasticity λ_k from Table 4, and the beta loadings of assets to factors (signs illustrated in Figure 4) to compute cross-asset elasticity. Each entry represents the price movement in bps of a row (column) asset, as induced by a $1 billion trading demand shock into a column (row) asset. As noted after Proposition 2, the model-implied cross-elasticity is symmetric, so we report only the upper half.

the traded FX factors. To induce the intermediaries to bear the incremental risk, prices of the traded FX factors change, which in turn affect the price of all other asset classes. Following Proposition 2, we combine the intermediary elasticity for traded FX factors with asset classes' return loadings on these factors to arrive at the own- and cross-elasticity between six asset classes (Table 8). Similar to Table 6, each number in Table 8 represents a row (column) asset's price movement in bps that is induced by a $1 billion trading demand shock into a column (row) asset, as channeled through both assets' exposure to the three traded FX factors.

Looking at the last column of Table 8, we note that although our exercise only makes use of asset loadings and intermediary elasticity, we recover two salient features of the US Treasury bonds (Treasurys). First, the price response to one's own trading demand shock is the smallest for Treasurys, corroborating the observation that the market for Treasurys is deep and liquid. Second, Treasurys have uniquely negative cross-elasticities with most other asset classes. The fact that a $1 billion trading demand shock to Treasurys raises its own price but depresses the price of other assets is reminiscent of Treasurys' "safe haven" property. Our estimation mimics this property because Treasurys load uniquely negatively on the commonly priced Carry factor.

We raise two cautions in interpreting our estimated cross-asset elasticity. First, our estimates capture only the elasticity due to exposure to the three traded FX factors. Our estimates may not, therefore, represent the total price response to a $1 trading demand shock into an asset, as these assets may also load on other risk factors that we do not capture. Second, our analysis implicitly assumes that the marginal intermediaries are the same across different markets. Nevertheless, our analysis highlights that if intermediaries are the marginal pricers in many markets and have limited risk-bearing capacity for commonly priced traded risk factors, then trading demand shocks affect asset prices across markets.

7 Conclusion

In conclusion, this paper studies the limited risk-bearing capacity of intermediaries and its implications for asset prices. To quantify intermediaries' risk-bearing capacities, we answer three questions: first, what risk factors are traded and therefore likely affected by trading demand shocks? Second, what is the "intermediary elasticity", or the price sensitivity to a marginal increase in risk induced by trading demand shocks? Third, what is the currency-level cross-elasticity arising from exposures to the traded FX factors? To answer these questions, we first extend priced non-diversifiable risks (Ross, 1976) to the representative intermediary (He and Krishnamurthy, 2017) and derive an original method to identify the most important traded risk factors: the Dollar, the Carry, and the Euro-Yen, which jointly account for 90% of the non-diversifiable risks borne by intermediaries when accommodating FX trading. We then use instrumental variables to estimate the intermediary elasticity of these traded risk factors, concluding that FX intermediaries have rather limited risk-bearing capacity. Finally, linking trading quantities to asset prices (Froot and Ramadorai, 2008; Koijen and Yogo, 2019) through *risks*, we use common factor exposure to derive novel cross-asset elasticity for a panel of 16 currencies and across 7 major asset classes, underscoring intermediaries' role in driving cross-asset dynamics (Haddad and Muir, 2021).

A distinguishing feature of our paper is the use of factor-level intermediary elasticity to inform cross-elasticity at both the currency and asset-class levels. At the heart of this cross-asset demand transmission are three elements: how trading demand shocks alter the amount of non-diversifiable factor risks borne by intermediaries, how risk prices change to induce intermediaries to absorb the incremental risks, and how different assets are exposed to these risks. Combining these three elements generates novel transmission patterns, where trading demand shocks in one market could affect prices in other markets by differential magnitudes and even directions. As intermediaries function at the juncture of financial markets and real sectors, how intermediaries transmit trading demand shocks to asset prices across different markets is essential to understanding intermediaries' role in the broader economy.

Bibliography

Adrian, T., E. Etula, and T. Muir. 2014. Financial intermediaries and the cross-section of asset returns. *The Journal of Finance* 69:2557–96.

Albuquerque, R., J. M. Cardoso-Costa, and J. A. Faias. 2024. Price elasticity of demand and risk-bearing capacity in sovereign bond auctions. *The Review of Financial Studies* hhae027.

An, Y. 2023. Flow-based arbitrage pricing theory. Working paper, Johns Hopkins University.

An, Y., Y. Su, and C. Wang. 2024. Quantity, risk, and return. Working paper, Johns Hopkins University.

Andrews, I., J. H. Stock, and L. Sun. 2019. Weak instruments in iv regression: Theory and practice. *Annual Review of Economics* .

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

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.

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

Chaboud, A., D. Rime, and V. Sushko. 2023. The foreign exchange market. Working Paper 1094, Bank for International Settlements.

Chaudhary, M., Z. Fu, and J. Li. 2023. Corporate bond multipliers: Substitutes matter. *Available at SSRN* .

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.

Coval, J., and E. Stafford. 2007. Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics* 86:479–512.

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

Du, W., B. Hébert, and A. W. Huber. 2022. Are intermediary constraints priced? *Review of Financial Studies* .

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

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

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

Duffie, D., M. J. Fleming, F. M. Keane, C. Nelson, O. Shachar, and P. Van Tassel. 2023. Dealer capacity and us treasury market functionality. *FRB of New York Staff Report* .

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

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. 2023. 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.

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

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

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.

Huber, A. W. 2023. Market power in wholesale funding: A structural perspective from the triparty repo market. *Journal of Financial Economics* 149:235–59.

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. *Journal of Finance* Forthcoming.

Jolliffe, I. 1986. *Principal component analysis.* Springer Series in Statistics.

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

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

Li, J., and Z. Lin. 2022. Price multipliers are larger at more aggregate levels. *Available at SSRN 4038664* .

Liao, G., and T. Zhang. 2020. The hedging channel of exchange rate determination. *International finance discussion paper* .

Lou, D. 2012. A flow-based explanation for return predictability. *Review of Financial Studies*

25:3457–89.

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.

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

Ready, R., N. Roussanov, and C. Ward. 2017. Commodity trade and the carry trade: A tale of two countries. *The Journal of Finance* 72:2629–84.

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

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

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

Siriwardane, E., A. Sundaram, and J. Wallen. 2022. Segmented arbitrage. Working Paper.

Tobin, J. 1978. A proposal for international monetary reform. *Eastern Economic Journal* 4:153–9.

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.

Wittwer, M., and J. Allen. 2024. Market power and capital constraints. Working Paper, Bank of Canada Staff Working Paper.

A Proofs

This appendix provides proofs omitted in the main text.

A.1 Solution for Traded Risk Factors

In this appendix, we present solutions for traded risk factors in Section 2.2.

We conduct Cholesky decomposition of $\text{var}(\mathbf{r})$ as $\mathbf{U}^\top\mathbf{U}$. Then, we define $\mathbf{g}_k = \mathbf{U}\mathbf{b}_k$ for each factor k. Equation (1) implies that the factor-level flow is

$$q_k = (\mathbf{b}_k^\top \text{var}(\mathbf{r})\mathbf{b}_k)^{-1}\mathbf{b}_k^\top \text{var}(\mathbf{r})\mathbf{f} = (\mathbf{g}_k^\top \mathbf{g}_k)^{-1}\mathbf{g}_k^\top \mathbf{U}\mathbf{f}. \tag{A1}$$

Moreover, the sequential optimization problem (3) becomes

$$\begin{aligned} &\max_{\mathbf{g}_k} (\mathbf{g}_k^\top \mathbf{g}_k)^{-1}\text{var}(\mathbf{g}_k^\top \mathbf{U}\mathbf{f}) \\ &\text{s.t. } \mathbf{g}_k^\top \mathbf{g}_j = 0 \text{ for } k \neq j. \end{aligned} \tag{A2}$$

This becomes a standard PCA problem that is solved by the eigenvalue decomposition of the matrix $\text{var}(\mathbf{U}\mathbf{f})$ (Jolliffe, 1986). The eigenvectors are $\mathbf{g}_k$ and the corresponding eigenvalues are proportional to the fraction of explained variance. Once we obtain $\mathbf{g}_k$, the portfolio weights are obtained by $\mathbf{b}_k = \mathbf{U}^{-1}\mathbf{g}_k$.

A.2 Invariance of Factors under Alternative Numeraire Currency

In this appendix, we prove that the factors constructed in Appendix A.1 remain unchanged when we alter the numeraire currency used to measure trading demand shocks and returns.

Suppose we switch from using USD to the N-th currency as the numeraire. We denote the trading demand shock from the N-th currency to the n-th currency as $\tilde{f}_n$ for $n = 1, 2, \ldots, N-1$, and the trading demand shock from the N-th currency to USD as $\tilde{f}_N$. Recall that f_n represents the trading demand shock from USD to the n-th currency. Because each trading demand shock f_n (for $n = 1, 2, \ldots, N-1$) can be broken down into a component from USD to the N-th currency and another from the N-th currency to the n-th currency, we can express this transformation as follows:

$$\tilde{\mathbf{f}} = (\tilde{f}_1, \tilde{f}_2, \ldots, \tilde{f}_{N-1}, \tilde{f}_N)^\top = \left(f_1, f_2, \ldots, f_{N-1}, -\sum_{n=1}^{N} f_n \right)^\top = \mathbf{C}\mathbf{f}, \tag{A3}$$

where we define the matrix

$$\mathbf{C} := \begin{pmatrix} 1 & 0 & \ldots & 0 & 0 \\ 0 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \ldots & 1 & 0 \\ -1 & -1 & \ldots & -1 & -1 \end{pmatrix}. \tag{A4}$$

Similarly, returns are now measured relative to the N-th currency. Specifically, $\tilde{r}_n$ for $n = 1, 2, \ldots, N-1$ represents the return from borrowing at the N-th currency's riskfree rate to invest in the n-th currency's riskfree rate. Similarly, $\tilde{r}_N$ denotes the return from borrowing at the N-th currency's riskfree rate to invest in the USD riskfree rate. The transformation of returns can thus be described as follows:

$$\tilde{\mathbf{r}} = (\tilde{r}_1, \tilde{r}_2, \ldots, \tilde{r}_{N-1}, \tilde{r}_N)^\top = (r_1 - r_N, r_2 - r_N, \ldots, r_{N-1} - r_N, -r_N)^\top = \mathbf{C}^\top \mathbf{r}. \tag{A5}$$

Now, we apply Appendix A.1 to analyze the factors using $\tilde{\mathbf{r}}$ and $\tilde{\mathbf{f}}$. Specifically, the variance of $\tilde{\mathbf{r}}$, given by $\text{var}(\tilde{\mathbf{r}}) = \mathbf{C}^\top \text{var}(\mathbf{r})\mathbf{C}$, can be decomposed as $\mathbf{C}^\top\mathbf{U}^\top\mathbf{U}\mathbf{C} = \tilde{\mathbf{U}}^\top\tilde{\mathbf{U}}$, where $\tilde{\mathbf{U}} := \mathbf{U}\mathbf{C}$. Subsequently, the eigenvalue decomposition is transformed to

$$\tilde{\mathbf{U}}\text{var}(\tilde{\mathbf{f}})\tilde{\mathbf{U}}^\top = \mathbf{U}\mathbf{C}\mathbf{C}\text{var}(\mathbf{f})\mathbf{C}^\top\mathbf{C}^\top\mathbf{U}^\top = \mathbf{U}\text{var}(\mathbf{f})\mathbf{U}^\top, \tag{A6}$$

where we use the fact that $\mathbf{C}\mathbf{C} = \mathbf{I}_N$. This derivation reveals that the eigenvectors $\mathbf{g}_k$ and eigenvalues are invariant. The resulting portfolio weights under the new numeraire currency are given by $\tilde{\mathbf{b}}_k = \tilde{\mathbf{U}}^{-1}\mathbf{g}_k = \mathbf{C}^{-1}\mathbf{U}^{-1}\mathbf{g}_k = \mathbf{C}^{-1}\mathbf{b}_k$. Hence, the factor returns also remain invariant, because $\tilde{\mathbf{b}}_k^\top\tilde{\mathbf{r}} = \mathbf{b}_k^\top(\mathbf{C}^{-1})^\top\mathbf{C}^\top\mathbf{r} = \mathbf{b}_k^\top\mathbf{r}$.

A.3 Proof of Proposition 1

Simplifying equation (4), we have

$$\begin{aligned}
&\mathbb{E}\left[-\exp\left(-\sum_{k=1}^{K}\gamma_k(y_k\mathbf{b}_k^\top\mathbf{R} - y_k R_F(1+\Delta p_k)\right)\right] \\
&= -\exp\left[-\sum_{k=1}^{K}\left(\gamma_k y_k\mathbb{E}[\mathbf{b}_k^\top\mathbf{R}] - \gamma_k R_F y_k(1+\Delta p_k) - \gamma_k^2 y_k^2\text{var}(\mathbf{b}_k^\top\mathbf{R})/2\right)\right],
\end{aligned} \tag{A7}$$

where the last equality uses the fact that $\text{cov}(\mathbf{b}_k^\top\mathbf{R}, \mathbf{b}_j^\top\mathbf{R}) = 0$ for $k \neq j$. Taking the first-order condition against y_k, we obtain

$$0 = \gamma_k\mathbb{E}[\mathbf{b}_k^\top\mathbf{R}] - \gamma_k R_F(1+\Delta p_k) - \gamma_k^2 y_k\text{var}(\mathbf{b}_k^\top\mathbf{R}). \tag{A8}$$

Because the optimal $y_k = -\hat{q}_k/\mu$, we obtain

$$1 + \Delta p_k = \frac{\text{var}(\mathbf{b}_k^\top\mathbf{R})\gamma_k\hat{q}_k/\mu + \mathbb{E}[\mathbf{b}_k^\top\mathbf{R}]}{R_F}. \tag{A9}$$

Specifically, when $\hat{q}_k = 0$, we have

$$1 = \frac{\mathbb{E}[\mathbf{b}_k^\top \mathbf{R}]}{R_F}. \tag{A10}$$

Taking the difference and using the assumption that $\text{var}(\mathbf{r}) = \text{var}(\mathbf{R})$, we obtain equation (5).

A.4 Proof of Proposition 2

Because factors have uncorrelated returns by equation (3), we can project the return of any currency n onto the factors and obtain

$$r_n = \sum_{k=1}^{K} \beta_{n,k} \mathbf{b}_k^\top \mathbf{r} + e_n, \tag{A11}$$

where e_n is the idiosyncratic return of currency n that is uncorrelated with any factor $\mathbf{b}_k^\top \mathbf{r}$. Hence, by the law of one price and equation (5), the price impact of currency n is

$$\Delta p_n = \sum_{k=1}^{K} \beta_{n,k} \Delta p_k = \sum_{k=1}^{K} \lambda_k \hat{q}_k \text{var}(\mathbf{b}_k^\top \mathbf{r}) \beta_{n,k}. \tag{A12}$$

Therefore, we have

$$\frac{\partial \Delta p_n}{\partial \hat{q}_k} = \frac{\partial \Delta p_k}{\partial \hat{q}_k} \times \frac{\partial \Delta p_n}{\partial \Delta p_k} = \lambda_k \text{var}(\mathbf{b}_k^\top \mathbf{r}) \times \beta_{n,k}. \tag{A13}$$

Next, equation (1) implies that $\partial \hat{q}_k / \partial \hat{f}_m = \beta_{m,k}$. Hence, we have

$$\frac{\partial \Delta p_n}{\partial \hat{f}_m} = \sum_{k=1}^{K} \frac{\partial \hat{q}_k}{\partial \hat{f}_m} \times \frac{\partial \Delta p_n}{\partial \hat{q}_k} = \sum_{k=1}^{K} \beta_{m,k} \times \lambda_k \text{var}(\mathbf{b}_k^\top \mathbf{r}) \times \beta_{n,k}, \tag{A14}$$

which is equation (8).

B 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 $100/(1+r^{JPY}_{1M})$ 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 consider trading flows in both the spot and the derivatives market.[45] 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 A1). Appendix Table A1 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

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

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

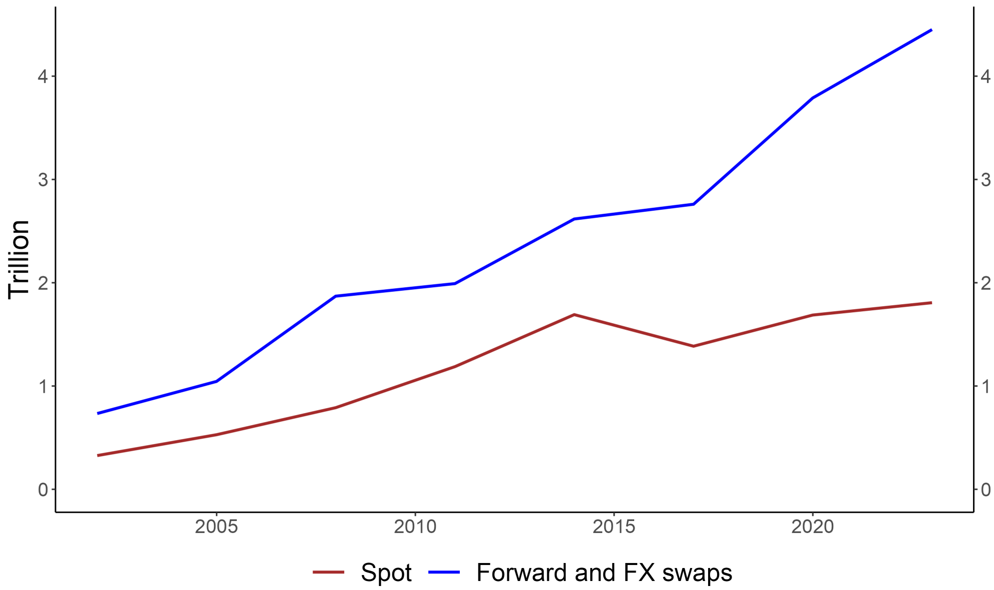

Notes: This figure plots the global daily volume of foreign exchange spot versus forward and FX swaps transactions involving USD. Daily volume is calculated as the average of all trading days in April of the survey year. The survey is conducted triennially from 2001 to 2022 by BIS.

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

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

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

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,

Table A2: **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.

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 elasticity estimation.

In Appendix Table A2, 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 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 elasticity to the risk factors because the spot market data alone may not provide an appropriate measure of the flow changes. Estimating elasticity 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 less price change per unit of additional risk.

C Additional Figures and Tables

Figure A3: **Cross-Asset Return Variation Outside of Crises**

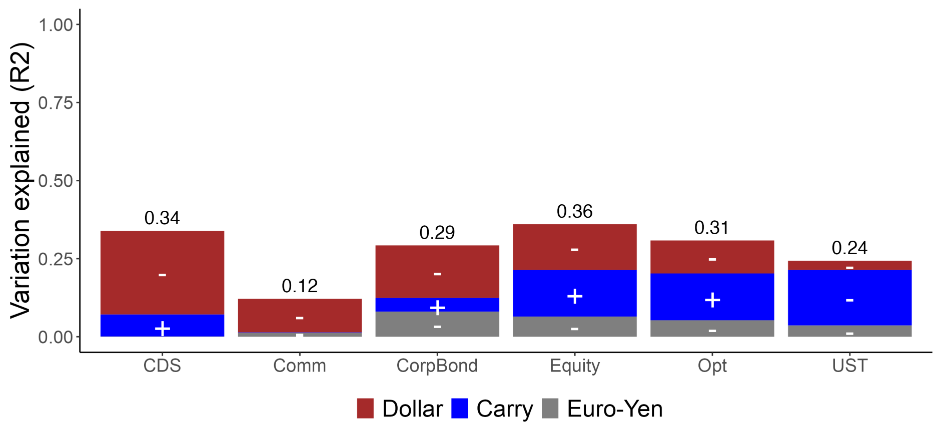

Notes: This figure decomposes the returns of individual assets into the Dollar, the Carry, and the Euro-Yen factors. The decomposition is achieved by regressing asset class monthly average excess return between 2000-02 and 2023-12 on returns from the three traded FX factors. We exclude the GFC (2007-07 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.

Table A3: **Top 3 PCs from FX Trading Flows**

Currency	PC 1	PC 2	PC 3
AUD	-0.03	0.03	0.12
CAD	-0.04	1.00	-0.06
CHF	-0.01	-0.02	-0.06
DKK	-0.00	-0.00	0.01
EUR	-1.00	-0.03	0.03
GBP	-0.02	-0.01	0.26
HKD	-0.00	-0.02	-0.00
ILS	-0.00	-0.01	-0.00
JPY	-0.04	-0.06	-0.95
KRW	-0.00	0.01	-0.00
MXN	-0.01	0.01	-0.00
NOK	0.00	0.01	0.01
NZD	-0.01	0.01	0.01
SEK	0.01	0.00	0.00
SGD	-0.01	-0.01	0.01
ZAR	-0.01	0.00	0.01
USD	1.17	-0.92	0.62
Flow Var explained	46%	21%	12%

Notes: This table presents the portfolio weights of the top 3 traded FX factors, constructed using a standard PCA of FX trading flows. We use weekly flow data for 16 non-USD currencies from September 2012 to December 2023. The portfolio weight of USD is computed as the negative sum of the weights of all other currencies.

Table A4: **Time-Varying λ for the Dollar Factor**

	Baseline (1)	Int. equity wealth (2)	CIP deviation (3)	Swap line usage (4)
Factor flow	0.107***	0.114***	0.160*	0.106***
	(0.037)	(0.033)	(0.092)	(0.029)
Flow × Int. wealth		-0.016		
		(0.031)		
Int. equity wealth		0.018		
		(0.130)		
Flow × CIP dev			0.063	
			(0.130)	
CIP deviation			0.172	
			(0.178)	
Flow × swap line				-0.056*
				(0.034)
Swap line usage				-0.449**
				(0.217)
Observations	386	383	386	385

Notes: This table reports the time-varying λ estimation results for the Dollar factor, as obtained using instrumental variables. "Int. equity wealth" is the cumulative (weekly) return of KBW Bank Index, demeaned and standardized. "CIP deviation" is measured by the weekly average AUD-JPY 3-month IBOR cross-currency basis, demeaned and standardized. "Swap line usage" is the week $t-1$ average amount outstanding at the Federal Reserve's central bank liquidity swap line, demeaned and standardized. 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.