University of Glasgow

Adam Smith Business School

WORKING PAPER SERIES

Foreign exchange order flow as a risk factor

Craig Burnside, Mario Cerrato and Zhekai Zhang

Paper No. 2023-03
January 2023

Foreign exchange order flow as a risk factor

Craig Burnside[*], Mario Cerrato[†], Zhekai Zhang[‡]

January 30, 2023

Abstract

We propose a novel pricing factor for currency returns motivated by the market-microstructure literature. Our factor aggregates order flow data to provide a measure of buying and selling pressure related to conventional currency trading strategies. It successfully prices the cross-section of currency returns sorted on the basis of interest rates and momentum. The association between our factor and currency returns differs according to the customer segment of the foreign exchange market. In particular, it appears that financial customers are risk takers in the market, while non-financial customers serve as liquidity providers.

Keywords: exchange rates, market microstructure, order flow, carry trade, currency momentum, crash risk, stochastic discount factor

JEL Codes: F31, G15.

[*]Duke University, University of Glasgow and NBER. E-mail: craig.burnside@duke.edu.

[†]Adam Smith Business School, University of Glasgow. E-mail: mario.cerrato@glasgow.ac.uk.

[‡]Institute of Economics and Finance, Nanjing Audit University. Email: zhangzhekai@nau.edu.cn

1 Introduction

Motivated by the market microstructure literature, we construct a reduced-form stochastic discount factor (SDF) model that explains the returns to popular currency trading strategies. Our model is based on customer order-flow data from one of the largest dealer banks in the foreign exchange market. The risk factor that we consider takes on more positive values when order flow, aggregated across currencies, indicates that customers are buying currencies favored by carry trade and momentum signals, and selling currencies disfavored by these signals. When this factor takes on more negative values, it indicates that customers are reversing or unwinding these trades. Our model successfully prices the cross-section of currency portfolios sorted by interest rates and momentum. This reflects the positive (negative) correlation between our factor and high (low) interest-rate currency portfolios, with the same pattern prevailing for strong (weak) momentum currency portfolios. A plausible interpretation of our findings, consistent with the market microstructure literature, is that when customers receive information that, in the aggregate, encourages them to take larger speculative positions in favor of both high interest-rate and strong momentum currencies, this causes these currencies to appreciate. When the arriving information causes them to reverse these positions, these currencies depreciate.

For those currencies with greater market share, we have access to order-flow classified according to customer type. These data indicate that the behavior of aggregate order flow is dominated by that of financial customers (hedge funds and asset managers), and that there are systematic differences in how the order flow of financial customers and nonfinancial customers (corporate and private clients) relate to currency returns. When the order flow of financial customers leans more towards taking carry trade or momentum positions, these investment strategies tend to do well. But we see the opposite pattern for nonfinancial customers. This suggests that order flow conveys different information to dealers depending on its origin within the customer base. It also suggests that a certain degree of risk sharing happens within the customer base, not just between customers and dealers (and between dealers). This may reflect that there are underlying shocks that drive exchange rate dynamics, and to which financial and nonfinancial customers have different ex-ante exposure.

Our work is closely related to the empirical microstructure literature that focuses on the relationship between order flow and bilateral exchange rates.[1] Like Lyons (2001) and

[1] See Sarno and Taylor (2001), Osler (2009), Evans (2011), Evans and Rime (2012), and Evans and Rime

Evans and Lyons (2002), we show that order flow can explain a large part of exchange rate variation, as well as, by extension, currency excess returns. They argue that this is because order flow maps a significant part of customers' private information into price discovery. The evidence we find is consistent with Evans and Lyons (2006), who show that there are significant differences, across customer segments, between the estimated price-impacts of order flow. Relatedly, Menkhoff et al. (2016) show that financial customer order flow contains information that has a long-term impact on currency returns, and that financial and nonfinancial customers trade in opposite directions, thus providing evidence of risk sharing taking place in the customer market. Ranaldo and Somogyi (2021) also document heterogeneity across customer segments.

Taking inspiration from this literature, we measure buying and selling pressure in the foreign exchange market, but do so by aggregating order flow (vis-a-vis the U.S. dollar) across currencies using standard trading signals for carry trade and momentum strategies. To illustrate, when we study the carry trade in isolation, our measure of aggregate order flow sums the value of buy orders for high interest rate currencies and the value of sell orders for low interest rate currencies, having normalized the measures of order flow to the scale of the market for each currency. When the value of this aggregate increases, we interpret this as customers, in general, favoring trades in the direction that carry-trade investing would predict. When the value of this factor decreases, we interpret this as customers reversing or unwinding these positions. We take a similar approach when studying momentum strategies in isolation. In this case, we aggregate the same order flow data but sign order flow based on a momentum signal. Then we combine our two measures by averaging them. This provides us a measure of the extent to which trading, in general, favors standard speculative strategies (or disfavors them).[2] We find that this single measure of aggregated order flow is able to price the cross-section of currency portfolios sorted on the basis of interest rates and momentum.

Our work is also closely related to the large empirical literature that uses SDF models to explain currency returns.[3] Villanueva (2007), Burnside et al. (2011), Burnside et al. (2011), and Burnside (2012) establish that traditional risk factors used to price equities do

(2019) for comprehensive overviews.

[2]Since the underlying order flow data are the same, if the carry trade and momentum signals agree at the level of an individual currency, then the contribution to both measures, and their average, is the same. If the signals disagree, then the contributions to the two measures are equal in magnitude but take the opposite signs, and so the contribution to the average is zero.

[3]Burnside (2012) and Lustig and Verdelhan (2012) review the early literature.

not correctly price carry trade and momentum portfolios. On the other hand, SDFs based on risk factors derived from currency-specific data have been reasonably successful in pricing portfolios of currencies. For example, Lustig et al. (2011) use a "high-minus-low" carry trade portfolio to price a set of currency portfolios sorted by interest rates. Menkhoff et al. (2012a) and Menkhoff et al. (2012b) use a measure of global currency volatility to price these same portfolios, as well as momentum portfolios. In our sample, our model appears to perform better than either of these traditional models in that it is able to price both the interest-rate-sorted and momentum-sorted portfolios, and our order-flow factor is robustly significant.

Another branch of the literature emphasizes currency crashes. Galati et al. (2007) find that excess returns to carry trades tend to reverse abruptly under market stress. They provide evidence from international banking data that currency flows are associated with these reversals. Brunnermeier et al. (2008) emphasize the role of risk averse market dealers who use the information in order flow to adjust the risk premium when they quote the spot rate. In their model, investors who engage in carry trades build their positions gradually, but liquidate their positions quickly, causing a currency crash. When market dealers anticipate a future unwinding by investors, they increase the risk premium associated with carry trade portfolios. Differently from Brunnermeier et al. (2008), in this paper we generalize that idea by extending it to the cross-section of currency returns, and we provide a natural empirical measure of trading pressure in the foreign exchange market. We find an association between signed order flow and currency returns that is broadly consistent with their notion of crash risk.

In Section 2, we describe the currency portfolios used in our empirical work. These include standard interest-rate-sorted and momentum-sorted portfolios used in the extant literature, as well as a set of portfolios sorted on the basis of order flow. In Section 3, we introduce our order-flow related pricing factors. Sections 4 and 5 contain the bulk of our empirical work, which is based on sample of weekly data from 2001 to 2012. We study the behavior of various currency portfolios in this period, as well as the performance of standard risk factors used in the prior literature. We then show cross-sectional asset pricing results for our order-flow based pricing factor. In Section 6, we discuss the behavior of order flow when it is disaggregated by customer segment. Section 7 concludes.

2 Currency Portfolios

Let $S_{k,t}$ be the exchange rate between the US dollar (USD) and foreign currency k, measured as foreign currency units (FCUs) per USD. Define $s_{k,t} = \ln S_t$. The logarithmic return to borrowing one USD in the short term money market and investing it in a short-term security denominated in foreign currency k, is

$$r_{k,t+1} = i^*_{k,t} - i_t - (s_{k,t+1} - s_{k,t}) \tag{1}$$

where i_t is the US interest rate and $i^*_{k,t}$ is the foreign interest rate. The uncovered interest parity (UIP) condition states that

$$E_t(s_{k,t+1} - s_{k,t}) = i^*_{k,t} - i_t, \tag{2}$$

or, equivalently, that

$$E_t r_{k,t+1} = 0, \tag{3}$$

where E_t is the expectations operator given information available at time t. That is, if the foreign interest rate exceeds the US interest rate, the foreign currency is expected to depreciate by the amount of the interest differential.

Let $F_{k,t}$ be the one period forward exchange rate between the same currencies, and let $f_{k,t} = \ln F_{k,t}$. Up to a log approximation, covered interest parity (CIP) implies that

$$i^*_{k,t} - i_t = f_{k,t} - s_{k,t}. \tag{4}$$

That is, the interest differential for currency k against the dollar is equal to currency k's forward discount.[4] Therefore, assuming that CIP holds, the log return to being long foreign currency k and short the USD is[5]

$$r_{k,t+1} = f_{k,t} - s_{k,t+1}. \tag{5}$$

Thus, under CIP, the UIP condition implies forward rate unbiasedness:

$$E_t s_{k,t+1} = f_{k,t}. \tag{6}$$

[4]Given that we quote exchange rates as FCUs per USD, $f_{k,t} - s_{k,t}$ measures how much cheaper it is to buy currency k forward rather than spot.

[5]Mancini-Griffoli and Ranaldo (2010) and Du et al. (2018) document that during and after the global financial crisis of 2008 there were substantial deviations from CIP as measured using money market interest rate data and bid-ask spreads. Prior to the financial crisis, these deviations were much less common and much smaller.

2.1 Carry Trade Portfolios

Carry trade strategies generally involve systematically managing a portfolio in which the investor borrows funds in low interest rate currencies and invests (or lends) in high interest rate currencies. Under uncovered interest parity, however, we would not expect this strategy to be profitable because $E_t r_{k,t+1} = 0$. However, the empirical failure of the UIP condition is well-documented.[6] In fact, it is widely understood that nominal exchange rates are well approximated, empirically, as random walks; i.e. $E_t s_{k,t+1} \approx s_{k,t}$.[7] When this is true

$$E_t r_{k,t+1} \approx i^*_{k,t} - i_t = f_{k,t} - s_{k,t}. \tag{7}$$

This fact provides motivation for carry trade strategies because it suggests that by systematically borrowing low interest rate currencies and lending in high interest rate currencies, the investor can expect to earn profits equal to the forward discount.

We base our empirical work on currency portfolios studied in the previous literature. Following Lustig et al. (2011), at each date t, we allocate the available currencies into five portfolios, labeled C1, C2, C3, C4 and C5, with C1 corresponding to the currencies with the lowest interest rates (equivalently, small values of the forward discount), and C5 containing those currencies with the highest interest rates (equivalently, large values of the forward discount). Each portfolio holds an equally weighted long position in its constituent currencies financed by borrowing dollars. Hence, the log return of the ith portfolio is

$$r^{\mathrm{C}i}_{t+1} = \sum_{k \in \mathcal{K}_{i,t}} \frac{1}{N_{i,t}} (f_{k,t} - s_{k,t+1}) = \sum_{k \in \mathcal{K}_{i,t}} \frac{1}{N_{i,t}} r_{k,t+1}, \tag{8}$$

where $\mathcal{K}_{i,t}$ is the set of currencies in the ith portfolio and $N_{i,t}$ is the number of currencies in the ith portfolio.

Lustig et al. (2011) use C1–C5 to construct two additional portfolios: the DOL portfolio and the HMLC portfolio. Their version of the DOL portfolio is an equally weighted average of the C1 through C5 portfolios. By contrast, we construct DOL as the equal weighted average of the currency excess returns for nine of the G10 currencies:[8]

$$r^{\mathrm{DOL}}_{t+1} = \frac{1}{9} \sum_{k \in \{\mathrm{G10}\}} r_{k,t+1}. \tag{9}$$

[6] Hansen and Hodrick (1980), Bilson (1981), Fama (1984) provide early tests. More recently, Engel (1996) and Burnside (2014) provide updated tests of UIP.

[7] The classic reference is Meese and Rogoff (1983).

[8] The nine currencies are the euro (EUR), Japanese yen (JPY), British pound (GBP), Swiss franc (CHF), Australian dollar (AUD), New Zealand dollar (NZD), Canadian dollar (CAD), Swedish krona (SEK), and Norwegian krone (NOK). The USD is left out as it is the base currency in our analysis.

This ensures that our DOL portfolio has a consistent definition across the different currency samples that we use by always measuring the tendency of the USD to depreciate or appreciate against the other G10 currencies.

The HMLC portfolio is a standard "high-minus-low" portfolio which takes a long position in the C5 portfolio and a short position in the C1 portfolio. In this sense, it can be thought of as a carry-trade portfolio that takes long positions in the highest interest rate currencies, financed by borrowing the lowest interest rate currencies. Its return is

$$r_{t+1}^{\mathrm{HMLC}} = r_{t+1}^{\mathrm{C5}} - r_{t+1}^{\mathrm{C1}}. \tag{10}$$

For the 2001–12 period, we form the C1–C5, HMLC and DOL portfolios using data for a set of 20 of the most liquid currencies according to trading volume.[9] The portfolios are formed on a weekly basis, each with a holding period of one week. Descriptive statistics are summarized in Table 1, with returns being expressed in percentage points per annum. Table 1 shows the mean return, standard deviation, skewness, kurtosis, Sharpe ratio, and the first order autocorrelation coefficient of the returns. We also report two coskewness measures relative to the returns to the DOL portfolio.[10] Portfolios with higher coskewness earn higher returns when global volatility is high. Thus, greater coskewness is often interpreted as making a portfolio more effective as a hedge against global volatility.

As Table 1 shows, the mean returns monotonically increase from portfolio C1 to portfolio C5 with the lowest return being 1.4% (on an annual basis) and the highest being 12.6%. The mean return of the DOL portfolio is 5.3%. This suggests that investors require a positive risk premium to invest in non-US short-term securities. Volatility also displays an increasing pattern moving from C1 to C5, but it does not rise in proportion to the expected return,

[9] The currencies in our data set are the EUR, JPY, GBP, CHF, AUD, NZD, CAD, SEK, NOK, MXN, BRL, ZAR, KRW, SGD, HKD, TRY, HUF, PLN, CZK, and SKK. We observe the exchange rates from the first week of November 2001 to the fourth week of March 2012. The appendix provides further details.

[10] Following Harvey and Siddique (2000) a direct measure for coskewness is

$$\beta_{\mathrm{SKS}} = \frac{E[\varepsilon_{t+1}\varepsilon_{M,t+1}^2]}{E[\varepsilon_{t+1}^2]^{1/2}E[\varepsilon_{M,t+1}^2]},$$

where ε_{t+1} is the innovation of the excess return of a portfolio, and $\varepsilon_{M,t+1}$ is the innovation of the excess return of some market factor (here we use the DOL factor). The innovations are constructed using first order autoregressive models for both the portfolio return and the DOL return.

The second coskewness measure is based on the regression

$$r_{t+1} = \beta_0 + \beta_1 r_{t+1}^{\mathrm{DOL}} + \beta_{\mathrm{SKD}}(r_{t+1}^{\mathrm{DOL}})^2 + u_{t+1},$$

where r_{t+1} is the return on some portfolio and $(r_{t+1}^{\mathrm{DOL}})^2$ is a proxy for market volatility.

so the Sharpe ratios also increase from C1 to C5. So high interest rate currencies still yield higher returns after a standard adjustment for risk. The HMLC portfolio (equivalent to C5-C1) has a large mean return (11.6%), and Sharpe ratio (0.99). The returns of all of the portfolios are negatively skewed, indicating the possibility of large negative realizations. However, for portfolio C1 the skewness coefficient is approximately zero, suggesting that it is less subject to large losses.

2.2 Momentum Portfolios

As documented by Burnside et al. (2011), Lustig et al. (2011), and Menkhoff et al. (2012b), momentum strategies in the foreign exchange market are also profitable. These strategies involve buying a basket of currencies with previously high returns and selling a basket with previously lower returns.

Similar to our approach for the carry trade, we form five momentum portfolios (M1, M2, M3, M4 and M5) based on the average value of each currency's log return over the previous four weeks. Portfolio M1 contains the currencies with the smallest lagged returns and portfolio M5 has the currencies with the largest lagged returns. In this regard, our portfolios are similar to those studied by Burnside et al. (2011), Menkhoff et al. (2012b), Fan et al. (2022), and Zhang (2022), which are formed on the basis of the previous month's return. However, in our case investors reshuffle their positions each week, rather than each month. We also consider a "high-minus-low" momentum portfolio that we denote as HMLM. Its return is equal to the return on the M5 portfolios minus the return on the M1 portfolio.

Table 1 provides a variety of summary statistics for these portfolios. The mean returns increase from portfolio M1 to portfolio M5 with the lowest return being 2.9% (on an annual basis) and the highest being 10.7%, although the pattern is not quite monotonic (portfolio M3 being the outlier). Consistent with the prior literature, we find that a strategy of holding the HMLM portfolio was profitable in historical data, with a large mean return (7.9%) and Sharpe ratio (0.75).

2.3 Order Flow and Exchange Rates

We also form portfolios based on order flow data. To do so, we use a unique data set, from one of the top foreign exchange dealers, covering more than eleven years (2001–2012) of weekly end-user order flow for up to 20 currencies.[11] Let $x_{k,t+1}$ denote the aggregate order flow (the total value of buy orders, net of sell orders) for currency k in the interval between periods t and $t+1$. Typically, empirical implementations of order flow models relate the change of the exchange rate to this flow, as well as to changes in observable fundamentals (such as the interest differential between the two currencies), and an error term. Our intention, here, is not to implement a specific order flow model. Instead, in our preliminary analysis, we demonstrate the apparent correlation between order flow and exchange rate changes at the weekly frequency.

In Table 2 we present estimates of the following equation for each currency:

$$s_{k,t+1} - s_{k,t} = a_k + x_{k,t+1} b_k + u_{k,t+1}, \tag{11}$$

where $u_{k,t+1}$ is an error term, and k indexes the currencies. Given that we measure exchange rates in FCUs per USD, and $x_{k,t+1}$ measures net buy orders of the foreign currency, we expect negative estimates of b_k. In fact, this is what we see in Table 2, with b_k being negative and statistically significant for 17 of our 20 currencies. This evidence is suggestive that order flow data may be useful in explaining exchange rate changes and the returns to currency investments. Order flow being significant at the weekly horizon mirrors the findings of Menkhoff et al. (2016), who demonstrate the predictive power of order flow over several days. It also reflects the results of several studies using longer-than-daily sampling frequencies that are surveyed by King et al. (2013).

2.4 Order Flow Portfolios

Order flow is not easily compared across currencies, due to heterogeneity in the volume of trade. To make such comparisons, we adjust currency k's order flow at time $t+1$ using the standard deviation of the order flow of currency k. To do this, we recursively define the sample variance of currency k's order flow as

$$\hat{\sigma}^2_{k,t} = \frac{1}{t} \sum_{s=1}^{t} (x_{k,s} - \bar{x}_{k,t})^2 \quad \text{with} \quad \bar{x}_{k,t} = \frac{1}{t} \sum_{s=1}^{t} x_{k,s}. \tag{12}$$

Then we define adjusted order flow as

$$y_{k,t+1} = \frac{x_{k,t+1}}{\hat{\sigma}_{k,t}}. \tag{13}$$

[11] The appendix provides further details of our data set.

We have found our results to be qualitatively robust to using a rolling-window definition of the standard deviation, as well as the full-sample standard deviation.[12]

At each week t, we sort the 20 currencies into five portfolios according to $y_{k,t}$, which are labeled O1, O2, ..., O5 where O1 consists of the currencies with greatest selling pressure (lowest, or most negative, order flow) and O5 consists of the currencies with the greatest buying pressure (most positive order flow). These are not tradable portfolios at time $t-1$ because the measure of order flow is contemporaneous to the return. Our purpose in studying these portfolios is, in fact, to measure the degree to which order flow and the returns are associated. We also define a buy-minus-sell (BMS) portfolio, which is long portfolio O5 and short portfolio O1.

Table 3(A) shows summary statistics for these portfolios. There is a clear monotonically increasing pattern in the average returns and Sharpe ratios across the O1–O5 portfolios. Unlike the interest rate sorted portfolios, C1–C5, the standard deviations of the returns do not vary much across the five portfolios. Unsurprisingly, the average of the O1–O5 portfolios (indicated by 'Avg' in Table 3) behaves similarly to the DOL portfolio in Table 1. The BMS portfolio earns a large positive average return, with a very large Sharpe ratio. These results, in a sense, confirm the notion that contemporaneous order flow is strongly positively correlated with exchange rate changes and currency returns.

Cerrato et al. (2011) argue that order flows from different segments of the customer market reflect the different information available to each segment, as well as their different motivations for trade. It is easy to imagine, for example, that leveraged hedge funds and corporate customers participate in the market for different reasons. To investigate whether there are systematic differences in the relationship between order flow and returns across customer-type we use data on order flow that are disaggregated into four categories: Asset Manager (AM), Hedge Fund (HF), Corporate (CO), and Private Client (PC). However, these data are only available for the nine G10 currencies, so we sort the currencies into four portfolios based on the magnitude of order flow in each of the customer segments. These results are reported in Table 3(B). For asset managers and hedge funds—referred to here as financial customers—the pattern across portfolios is the same as for aggregate order flow. The portfolios with the most buying pressure earn the largest returns. For corporate customers and private clients—referred to here as nonfinancial customers—the pattern is reversed,

[12]Duplicates of many of our tables using the full-sample standard deviation are available on request.

and sharply so for the latter category. The portfolios with the most buying pressure from nonfinancial customers earn negative returns, while the ones with the most selling pressure earn positive returns. Cerrato et al. (2011) show that these nonfinancial customers tend to act as liquidity providers. The evidence in Table 3(B) seems consistent with this view, in that currencies being bought by financial customers do better, while the opposite is true for nonfinancial customers.

Next, we compare the informational content of order flow with that of forward discounts and volatility innovations. Menkhoff et al. (2012a) show that a global volatility proxy contains important information which can be used to price returns of carry trade portfolios. Relatedly, Menkhoff et al. (2012b) show that momentum strategies are more profitable among currencies that have greater idiosyncratic volatility. In both cases, the implication is that volatility has an association with the riskiness of, and return to, holding different currencies and currency portfolios. We believe that the apparent importance of volatility is strongly linked to order flow and that, in fact, order flow contains the relevant information to price the returns of carry trade and momentum portfolios.

To provide the reader with a first intuitive view of this, we double-sort our currencies in two different ways with the results being shown in Tables 4 and 5. In Table 4, we first sort our currencies into three portfolios based on their forward discounts. Thereafter, within each portfolio, we sort currencies into two bins based on the magnitude of order flow.[13] The main conclusion of Table 4 is that even after conditioning on the forward discount (i.e., choosing a column in the table), buying a portfolio with the highest buying pressure (high order flow) and selling a portfolio with the highest selling pressure (low order flow), gives a positive and statistically significant return. In other words, taking forward discounts into account does not drive out order flow as an important apparent determinant of currency returns.

In Table 5, we first sort our currencies into three portfolios based on their idiosyncratic volatility innovation, and thereafter on the magnitude of order flow.[14] Again, even after conditioning on the idiosyncratic volatility innovations (i.e., choosing a column in the table), a portfolio of the currencies with the highest buying pressure has an economically and

[13]We build a total of just six portfolios due to the limited number of currencies in our sample.

[14]The idiosyncratic volatility innovations is measured in a similar fashion to how Menkhoff et al. (2012a) construct their risk factor, DVOL. For each week, and each currency, we average the absolute daily spot rate changes to proxy for the volatility of that currency in that week. We then model the volatility time series of each currency as an AR(1) process and take the residual term from the model as a proxy for the idiosyncratic volatility innovation of that currency.

statistically significantly higher return than the one with the greatest selling pressure.

3 Order-Flow Factors

The empirical results presented in Tables 3–5 suggest that order flow contains significant information that could be relevant for pricing the returns to currency portfolios. In this section, we propose novel pricing factors based on order flow that are motivated by microstructure models and the prevalence of carry and momentum trading in foreign exchange markets.

3.1 A Carry-Trade Order-Flow Factor

Our first factor is an aggregate order flow measure motivated by carry-trading considerations. This factor, which we denote as CO, is defined as

$$\mathrm{CO}_{t+1} = \frac{1}{N_t} \sum_{k \in N_t} y_{k,t+1} \cdot \mathrm{sign}(f_{kt} - s_{kt}), \tag{14}$$

where N_t denotes the total number of foreign currencies in the available data. If investors build portfolios based on carry-trade considerations, we might expect $y_{k,t+1}$ to be positive for currencies for which $f_{kt} - s_{kt}$ is positive, and negative for currencies for which $f_{kt} - s_{kt}$ is negative. Thus, we would expect CO to generally be positive.[15] But $y_{k,t+1}$ should also reflect news that arrives after investors form their portfolios, because it measures order flow between periods t and $t+1$. If arriving news is favorable to carry trades, we would expect CO to be especially high. On the other hand, if news arrives that induces investors to cash out their carry trade positions, CO will fall, and possibly even turn negative. In a sense, therefore, CO can be interpreted as a factor that measures the degree of sentiment in favor of carry trading as reflected in customer order flow.[16]

[15]This intuition is based on the notion that investors form carry-trade portfolios from the perspective of a US investor, going long (short) those currencies whose interest rates are higher (lower) than the US interest rate. If, instead, investors form these portfolios from a dollar-neutral perspective, the relevant signal variable might reflect the size of the interest differential relative to the median in the sample of currencies.

[16]Burnside (2012) suggests that a significant part of trading activity in foreign exchange markets is triggered by carry trade investors. Breedon et al. (2016) show that there is a strong relationship between order flow data and currency forward premia.

3.2 A Momentum Order-Flow Factor

Our second factor is an aggregate order flow measure motivated by momentum-trading considerations. This factor, which we denote as MO, is defined as

$$\mathrm{MO}_{t+1} = \frac{1}{N_t} \sum_{k \in N_t} y_{k,t+1} \cdot \mathrm{sign}(\bar{r}_{4t}), \tag{15}$$

where

$$\bar{r}^4_{kt} = r_{kt} + r_{kt-1} + r_{kt-2} + r_{kt-3}. \tag{16}$$

If investors build portfolios based on momentum considerations, we might expect $y_{k,t+1}$ to be positive for currencies for which $\bar{r}^4_{kt}$ is positive, and negative for currencies for which $\bar{r}^4_{kt}$ is negative. Thus, like CO, we would expect MO to generally be positive.[17] If arriving news is favorable to momentum trades, we would expect MO to be especially high. On the other hand, if news arrives that induces investors to cash out their momentum positions, MO will fall, and possibly even turn negative. In a sense, therefore, MO can be interpreted as a factor that measures the degree of sentiment in favor of momentum trading as reflected in customer order flow.

3.3 An Overall Currency Speculation Factor

Our third factor is motivated by the following discussion. Suppose that at time t, currency k has a higher interest rate than the USD, and has positive momentum. This means that $y_{k,t+1}$ is counted positively in the construction of CO and MO, and contributes the same amount to both measures. If we then observed that $y_{k,t+1}$ was large and positive, it would suggest that, for whatever reason, customers are increasing traditional speculative positions in currency k, in general. If we observed that $y_{k,t+1}$ was large and negative, it would suggest that customers are reversing traditional speculative positions in currency k, in general. A similar interpretation is possible when the interest rate and momentum signals are both negative.

If the carry and momentum signals take on different signs, however, it is difficult to assign the same meaning to the value of $y_{k,t+1}$. Suppose, for example, that, at time t, currency k has a high interest rate and has negative momentum. This means that $y_{k,t+1}$ is counted positively

[17] As for carry, this intuition is based on the notion that investors form momentum portfolios from the perspective of a US investor, going long (short) those currencies which have positive (negative) momentum versus the USD. If, instead, investors form these portfolios from a dollar-neutral perspective, the relevant signal variable might reflect the size of $\bar{r}^4_{kt}$ relative to the median in the sample of currencies.

in the construction of CO, but negatively in the construction of MO. If we then observed that $y_{k,t+1}$ was large and positive, it would suggest that, for whatever reason, customers are increasing carry positions in currency k, but decreasing momentum positions in currency k, given that momentum would suggest a negative value of $y_{k,t+1}$.

For this reason, we propose a third factor, denoted CM, which is the simple average of CO and MO:

$$\mathrm{CM}_t = \frac{\mathrm{CO}_t + \mathrm{MO}_t}{2}. \tag{17}$$

Suppose that for currency k, $\mathrm{sign}(f_{kt} - s_{kt}) = \mathrm{sign}(\bar{r}^4_{kt})$. Then $y_{k,t+1}$ contributes the equal amounts to CO, MO and CM. However, when $\mathrm{sign}(f_{kt} - s_{kt}) \neq \mathrm{sign}(\bar{r}^4_{kt})$, $y_{k,t+1}$ contributes opposite amounts to CO and MO and, therefore, nothing to CM. Therefore, CM measures signed order flow summed across all currencies for which the trading signals point in the same direction. Thus, when we observe a large positive value of CM it suggests that customers are increasing their traditional speculative currency positions, in general. When we observe a large negative value of CM it suggests that customers are decreasing their traditional speculative currency positions, in general.

4 The Risk Exposure of Currency Portfolios

In this section, we measure the risk exposures of the portfolios we constructed in Section 2.1. To do so, we follow the standard approach in the literature, which is to perform time series regressions of the returns of these portfolios on vectors of risk factors. These risk factors include ones selected from the literature, as well as the novel order-flow based factors we introduced in Section 3. Each time series regression is of the form

$$r^e_{i,t} = \alpha_i + z_t'\beta_i + \epsilon_{i,t}, \qquad i = 1, \ldots, N, \quad t = 1, \ldots, T, \tag{18}$$

where $r^e_{i,t}$ is the excess return of portfolio i at time t, z_t is an $n \times 1$ vector of risk factors, N is the number of portfolios and T is the sample size. In this part of our analysis, we consider the five interest rate-sorted currency portfolios, C1–C5, and the five momentum-sorted currency portfolios, M1–M5.

4.1 Betas of Traditional Pricing Factors

We begin by considering two risk factors similar to those proposed by Lustig et al. (2011): DOL and HMLC. Overall, the results, shown in Table 6, are in line with what has been documented in the empirical literature. The betas for DOL are scaled near unity, although they are somewhat smaller for C1, C5, M1 and M5 than they are for the other portfolios. The betas for the HMLC factor increase across the interest-rate sorted portfolios. With a beta of -0.26, C1 has a negative exposure to HMLC, indicating that it is a hedge against carry trade risk. By contrast, the beta is large and positive for C5 (0.74), indicating that it is highly exposed to carry trade risk. These results are not surprising given the construction of the factors.[18] The momentum portfolios have positive exposure to HMLC, across the board. The betas with respect to HMLC generally decrease from M1, which has by far the largest beta, to M5.

Table 7 shows results for factors similar to those used by Menkhoff et al. (2012a), which are DOL and a global volatility innovation factor (DVOL). The DVOL factor is measured as the cross-sectional average of the intra-week volatility innovation for each currency in our sample (see footnote 14). In Menkhoff et al. (2012a) the same measure is used but it is computed on an intra-month basis. Again, the results are in line with what has been documented in the literature. The pattern in the betas for DOL is similar to what we observed for the DOL-HMLC model. For DVOL, the betas are positive for the low-interest-rate currency portfolios (C1 and C2) and negative, and increasingly so, for the high-interest-rate portfolios (C3, C4 and C5). This indicates that when global currency volatility rises, high interest rate currencies tend to do poorly, while low interest rate currencies act as a hedge against increasing volatility. The momentum portfolios have small negative exposures to DVOL, across the board, although in many cases the estimated betas are not statistically significant.

4.2 Betas of the Order-Flow Factors

4.2.1 Carry

Table 8 shows results obtained using our aggregate carry-trade order-flow risk factor, CO, in tandem with the DOL factor. The pattern in the betas for DOL is similar to what we

[18]This follows from the fact that DOL is similar to the average of C1–C5 while HMLC is C5 minus C1. See Burnside (2010) for further details.

observed for the DOL-HMLC and DOL-DVOL models. The results indicate that portfolios with higher interest rates (C3, C4, and C5) have positive and statistically significant exposure to CO. The lower interest rate portfolios (C1 and C2) have negative and, in the case of C1, statistically significant exposure to CO. The betas for C3–C5 are positive and statistically significant. The betas are monotonically increasing as we move from C1 to C5. These results mean that when the order flow data suggest stronger trading pressure consistent with the carry trade, i.e. when CO increases, the high interest rate portfolios earn higher returns and the low interest rate portfolios earn lower returns. The pattern reverses if investors reverse their carry trade holdings and CO decreases. Consequently, low interest rate portfolios act as hedges against a reversal of investors' carry trade positions, while high interest rate portfolios are exposed to this risk. The momentum portfolios also have positive exposure to CO, and in many cases it is statistically significant. The pattern in the betas is generally increasing from M1 to M5, but M2 breaks the order by having the second largest estimated beta.[19]

4.2.2 Momentum

Table 9 shows results obtained using our momentum carry-trade order-flow risk factor, MO, in tandem with the DOL factor. The pattern in the betas for DOL is similar to what we observed for the DOL-CO model. The carry portfolios (C1–C5) have small and statistically insignificant exposures to MO. For the momentum portfolios, the betas are monotonically increasing as we move from M1 to M5. These results mean that when the order flow data suggest stronger buying pressure consistent with momentum trading, i.e. when MO increases, the currency portfolios with greater momentum earn higher returns and those with less momentum earn lower returns. Consequently, portfolios with little momentum act as hedges against a reversal of investors' momentum positions, while portfolios with the most momentum are exposed to this risk.

[19]Bussiere et al. (2018) and Burnside (2019) document structural breaks in standard UIP regressions (regressions of exchange rate changes on forward discounts) around the time of the Global Financial Crisis of 2008. In previous work, we investigated the stability of the betas for the HMLC, DVOL and CO factors across subsamples in our data set; in particular, before and after the Global Financial Crisis of 2008. Focusing on the C1–C5 portfolios, we found betas to be very stable for the DOL-HMLC model, but this is not surprising given the construction of DOL and HMLC. On the other hand, there was a fairly striking decrease in the magnitude and statistical significance of the exposures of the different portfolios to DVOL and an apparent increase in the importance of DOL within the DOL-DVOL model. For the DOL-CO model the main difference across subsamples was an increase of the exposures of the higher interest rate currency portfolios (C3–C5) to both factors. These results are available upon request.

4.2.3 Currency Speculation

Table 10 shows results obtained using our overall currency speculation order-flow risk factor, CM, in tandem with the DOL factor. The pattern in the betas for DOL is similar to what we observed for the DOL-CO and DOL-MO models. The betas increase monotonically across the carry portfolios (C1–C5) and the momentum portfolios (M1–M5). These results mean that when the order flow data suggest stronger buying pressure consistent with standard currency speculation strategies, i.e. when CM increases, the currency portfolios with greater carry or momentum earn higher returns and those with less earn lower returns. Consequently, portfolios with little carry or momentum act as hedges against a reversal of investors' speculative positions, while portfolios with the most carry or momentum are exposed to this risk.

Figure 1 summarizes the relationship between the betas that are unique to each of the five models and the average returns of the ten test portfolios. While the betas with respect to HMLC and DVOL line up reasonably with the expected returns of the carry portfolios, they capture little of the variation in expected return across the momentum portfolios. For our MO factor, the betas line up quite well with the average returns of the momentum portfolios, but not with those of the carry portfolios. However, the betas for our CO factor, and, more especially, our CM factor, line up quite well with the average returns of all of the portfolios.

5 Cross-Sectional Asset Pricing

In this section, we use a generalized method of moments [GMM, Hansen (1982)] approach to estimate linear stochastic discount factor (SDF) models, discussed in Cochrane (2009), and used by Lustig et al. (2011), Burnside et al. (2011), and Menkhoff et al. (2012a) among many others. Let r^e be an $N \times 1$ vector of excess returns where N is the number of test assets. If m_t is an SDF for these returns, then

$$E(r^e m) = 0 \tag{19}$$

where E is the unconditional expectations operator. As is standard in the literature, we specify the SDF as a linear function of a $n \times 1$ vector of risk factors, z:

$$m = 1 - (z - \mu)'b, \tag{20}$$

where $\mu = E(z)$ and b is a $n \times 1$ vector of parameters. Given this definition, the mean of the SDF is normalized to 1.

When equation (19) is combined with equation (20) it becomes

$$E(r^e) = \text{cov}(r^e, z)b. \tag{21}$$

Our other moment restriction is

$$E(z) = \mu. \tag{22}$$

This motivates the use of the following GMM estimators for b and μ

$$\hat{b} = (C'WC)^{-1}C'W\bar{r}^e, \tag{23}$$

$$\hat{\mu} = \bar{z}, \tag{24}$$

where $\bar{r}^e$ is the sample mean of r^e, $\bar{z}$ is the sample mean of z, C is the sample covariance matrix between r^e and z, and W is some positive definite weighting matrix. For the results reported in this paper, we set $W = I_N$.[20]

Letting $\Sigma_z = E[(z-\mu)(z-\mu)']$, equation (21) can also be written as

$$E(r^e) = \left[\text{cov}(r^e, z)\Sigma_z^{-1}\right](\Sigma_z b) = \beta\lambda, \tag{25}$$

with $\beta = \text{cov}(r^e, z)\Sigma_z^{-1}$ being an $N \times n$ matrix of factor betas, and $\lambda = \Sigma_z b$ being a $n \times 1$ vector of risk prices. This is the beta representation of the pricing model, which we also estimate using GMM, as described in Cochrane (2009), and in the appendix to Burnside (2011). Our estimation procedure is equivalent to Fama and MacBeth (1973)'s method, with standard errors being calculated as per Shanken and Zhou (2007).

When estimating either the SDF representation of the model or the beta representation, it is important that the matrix $\text{cov}(r^e, z)$ has full column rank (i.e. its rank should be n). When this condition fails, the model is not properly identified, both estimators have non-standard asymptotic distributions, and tests for the validity of the model also have non-standard distributions as discussed in Burnside (2016). Therefore, we perform the tests proposed by Kleibergen and Paap (2006) (KP) for testing the rank of $\text{cov}(r^e, z)$. We mainly work with models where $n = 2$. If $\text{cov}(r^e, z)$ has rank 0, it means neither risk factor is correlated with the return vector; i.e. all the elements of $\text{cov}(r^e, z)$ are 0. If $\text{cov}(r^e, z)$ has rank 1, it means

[20]Details of the computation of the parameter estimates and standard errors are provided in the online appendix to Burnside (2011).

one risk factor is uncorrelated with the return vector or a linear combination of the two risk factors is uncorrelated with the return vector; i.e. one of the columns of $\text{cov}(r^e, z)$ is zero or the two columns of $\text{cov}(r^e, z)$ are proportional to one another.

As test assets, we use the returns to the ten portfolios described above (C1–C5 and M1–M5). In the tables that follow, we report parameter estimates, and standard errors. When estimating the SDF representation, we report Hansen and Jagannathan (1997)'s distance measure as a test of the model's fit. When estimating the beta representation, we report the results of a test for whether the pricing errors are zero.

5.1 The DOL-HMLC Model

In Table 11, we start with the DOL and HMLC factors proposed by Lustig et al. (2011). Qualitatively, the results are in line with what has been documented in the existing literature. The SDF parameter (b) for the HMLC factor is positive, although only statistically significant at around the 12% level. The associated risk price (λ) is positive and statistically significant at the 7% or 1% level depending on the procedure used to compute standard errors. For the DOL factor both parameters are positive, but neither is statistically significant.

The cross-sectional fit of the model is modest, with an $R^2 = 0.39$. The modest fit of the model can also be seen in Figure 2 which plots the model-predicted expected returns, $C\hat{b}$, against the average returns of the ten test assets, $\bar{r}^e$. While the model does quite well in fitting the carry portfolios (C1–C5), it does rather poorly in explaining the momentum portfolios (M1–M5). In fact, it predicts that the M5 portfolio should have a lower expected return than the M1 portfolio. This is not puzzling: As we saw above, the M1 portfolio has a much larger exposure to HMLC risk than the M5 portfolio does.

While the model passes the Hansen-Jagannathan specification test (p-value of 0.16), it is rejected based on the Fama-MacBeth specification test (p-value of 0.01). The KP test strongly rejects the null of reduced rank.

5.2 The DOL-DVOL Model

Table 12 shows results for a model similar to the one used by Menkhoff et al. (2012a), which includes DOL and DVOL as factors. Qualitatively, the results are in line with what has been documented in the existing literature. The SDF parameter and the risk price of DVOL are both negative, indicating that portfolios with greater exposure to higher volatility (i.e.

lower returns when volatility increases) have higher mean returns. However, neither $\hat{b}_{\text{DVOL}}$ nor $\hat{\lambda}_{\text{DVOL}}$ is statistically significant at conventional significance levels, except when we use the Fama-MacBeth method.

The cross-sectional fit of the model is slightly better than that of the DOL-HMLC model, because the model fits the momentum portfolios somewhat better, and the carry portfolios almost as well. This is illustrated in Figure 3.

While the model passes the Hansen-Jagannathan specification test (p-value of 0.30), it is rejected based on the Fama-MacBeth specification test (p-value of 0.02). On the other hand, the KP test only rejects the null hypothesis of reduced rank at the 24% level. This may reflect the degree of imprecision with which the betas are estimated for the DVOL factor.

5.3 The DOL-CM Model

Rather than working with the three different order-flow related models from Section 4, here we focus on a model that uses the DOL factor in tandem with the overall currency speculation order-flow factor, CM. Table 13 shows cross-sectional asset pricing results for this model. The empirical evidence in Table 13 strongly supports CM as a pricing factor. The SDF parameter (b) and risk price (λ) for the CM factor are positive and statistically significant. Thus, portfolios with more positive exposure to CM carry larger risk premia. The cross-sectional fit of the model is good, with $R^2 = 0.70$. The fit of the model is illustrated in Figure 4. Comparing this graph to Figures 2 (DOL-HMLC) and 3 (DOL-DVOL), we can see that model does a little worse at fitting the interest-rate sorted portfolios (C1–C5) than the DOL-HMLC model, but it does much better than both models in fitting the momentum sorted portfolios (M1–M5). The overall degree of fit is substantially better.

The model is not rejected at the conventional significance levels based on the HJ distance measure and the Fama-MacBeth pricing error test. The KP test strongly rejects the null hypothesis of reduced rank at less than the 1% level.

6 Disaggregated Order Flow

As discussed above, we have data on order flow that are disaggregated by customer segments. As we saw in Section 2, there are differences in how order-flow is contemporaneously related to currency returns across customer segments: Portfolios of currencies with more buying pressure from financial customers have higher returns than those with less buying pressure.

But the reverse pattern is observed for nonfinancial customers. This suggests that different customer segments act in systematically different ways. This might reflect that they have access to different information, or have different motives for trade, possibly rooted in different *ex-ante* exposure to risk.

To explore further, we construct alternative order-flow factors corresponding to the Asset Manager, Hedge Fund, Corporate, and Private Client customer segments. These factors are conceptually the same as the CO, MO, and CM factors, but they measure whether each customer type is trading in ways that are consistent with carry trade signals, momentum signals, or standard currency speculation signals. Consequently, we denote these factors as COAM, COHF, COCO, COPC (related to CO), MOAM, MOHF, MOCO, MOPC (related to MO), and CMAM, CMHF, CMCO and CMPC (related to CM). To take Asset Managers as an example, for each currency we define normalized order flow as

$$y_{k,t+1}^{\mathrm{AM}} = \frac{x_{k,t+1}^{\mathrm{AM}}}{\hat{\sigma}_{k,t}}, \tag{26}$$

where $x_{k,t+1}^{\mathrm{AM}}$ is raw order flow for the Asset Manager segment, and $\hat{\sigma}_{k,t}$ is the recursively-defined standard deviation of order flow summed across all customer segments. Given this definition our normalized order-flow measures by customer segment aggregate in the same way as the raw order flow data:

$$y_{k,t+1}^{\mathrm{AM}} + y_{k,t+1}^{\mathrm{HF}} + y_{k,t+1}^{\mathrm{CO}} + y_{k,t+1}^{\mathrm{CP}} = y_{k,t+1} \tag{27}$$

The COAM factor is defined as

$$\mathrm{COAM}_{t+1} = \frac{1}{N_t} \sum_{k \in N_t} y_{k,t+1}^{\mathrm{AM}} \cdot \mathrm{sign}(f_{kt} - s_{kt}), \tag{28}$$

The other customer-segment factors are defined similarly.

Table 14(a) shows the correlation matrix between the disaggregated COxx factors and the CO factor, redefined using only the nine major currencies. It shows that the signed order flow of financial customers, as measured by COAM and COHF, is quite highly correlated with overall signed order flow, as measured by CO. However, COAM and COHF are not highly correlated with each other. The order flow of nonfinancial customers, as measured by COCO and, especially, COPC, has much lower correlation with CO. We even see negative correlation between COPC and both COAM and COHF. This suggests systematic trading differences across customer types.

Table 14(b) shows the correlation matrix between the disaggregated MOxx factors and the MO factor, redefined using only the nine major currencies. It shows that the signed order flow of financial customers, as measured by MOAM and MOHF, is highly correlated with overall signed order flow, as measured by MO. However, MOAM and MOHF are not highly correlated with each other. As with the carry trade factors, the order flow of nonfinancial customers, as measured by MOCO and MOPC, has much lower correlation with MO. Furthermore, both MOCO and MOPC are negatively correlated with MOAM and MOHF. Once again, this suggests systematic trading differences across customer types.

Table 14(c) shows the correlation matrix between the disaggregated CMxx factors and the CM factor, redefined using only the nine major currencies. Unsurprisingly, it shows that the signed order flow of financial customers, as measured by CMAM and CMHF, is quite highly correlated with overall signed order flow, as measured by CM. However, CMAM and CMHF are not highly correlated with each other. The order flow of nonfinancial customers, as measured by CMCO and CMPC, has much lower correlation with CM. Furthermore, both CMCO and CMPC are negatively correlated with CMAM and CMHF.

These findings suggest a risk-sharing story consistent with the one in Menkhoff et al. (2016), in that different group of customers (i.e. financial and non-financial) appear to trade in different directions and, therefore, risk sharing takes place in the customer market, not just in the inter-dealer market as emphasized in the early literature on order flow in currency markets.[21] That said, order flow is dominated by that of financial customers. For example, if we consider our CMxx factors, the financial order flow variables CMAM and CMHF account for 95% of the variation in CM (when it is redefined using only the nine major currencies).

Further evidence of systematic differences in behavior across customer segments is found in Figure 5. There we plot the betas of the disaggregated overall order-flow factors (i.e. the CMxx factors). In each case, a two factor model, with DOL as the other factor, is estimated. It is notable that the pattern in the betas for financial customers much more closely matches the pattern we observed for the CM factor. For both carry and momentum portfolios, the pattern in the betas is generally increasing as we move from low return to high return portfolios. When we move to Corporate customers, the pattern is reversed for the momentum portfolios. Both patterns are reversed for private clients.

[21]For the equity market, Barber and Odean (2013) show that private investors (i.e. uninformed investors) tend to lose money from trading.

7 Conclusion

We have demonstrated that, at the weekly frequency, order-flow is closely associated with systematic patterns in currency returns. We have shown that if currencies are sorted on the basis of aggregated normalized order-flow, portfolios of currencies with stronger buying pressure tend to appreciate relative to currencies with weaker buying (or strong selling) pressure. At the disaggregated level, we see the same pattern when we use the order-flow of financial customers (hedge funds and asset managers). However, the pattern is reversed when we use the order-flow of non-financial customers (corporates and private customers). This suggests that a form of risk sharing takes place in the foreign exchange market, not just between dealers and non-dealers, but within the confines of the non-dealer customer base.

We have also used order-flow based risk factors in a traditional SDF approach to cross-sectional asset pricing. The novel risk factor that we considered increases when customers buy currencies that are favored by both carry trade and momentum signals, and sell currencies disfavored by both of these signals. We argued that this risk factor acts as a measure of customers' willingness to engage in standard currency speculation strategies. The model that we estimated using this factor is successful in pricing the cross-section of interest-rate and momentum-sorted currency portolfios in our data set. In this respect, the model appears more successful than other notable reduced-form models from the literature.

References

Barber, B. M., and T. Odean. 2013. The Behavior of Individual Investors. In G. M. Constantinides, M. Harris, and R. M. Stulz (eds.), *Handbook of the Economics of Finance*, vol. 2B, chap. 22, pp. 1533–1570. Elsevier.

Bilson, J. F. O. 1981. The "Speculative Efficiency" Hypothesis. *The Journal of Business* 54:435–451.

Breedon, F., D. Rime, and P. Vitale. 2016. Carry trades, order flow, and the forward bias puzzle. *Journal of Money, Credit and Banking* 48:1113–1134.

Brunnermeier, M. K., S. Nagel, and L. H. Pedersen. 2008. Carry trades and currency crashes. *NBER Macroeconomics Annual* 23:313–348.

Burnside, C. 2010. A Note on Sorts and Cross-Sectional Regressions. Duke University.

Burnside, C. 2011. The cross section of foreign currency risk premia and consumption growth risk: Comment. *American Economic Review* 101:3456–76.

Burnside, C. 2012. Carry Trades and Risk. In J. James, I. Marsh, and L. Sarno (eds.), *Handbook of Exchange Rates*, Wiley Handbooks in Financial Engineering and Econometrics, chap. 10, pp. 283–312. Wiley.

Burnside, C. 2014. The Carry Trade in Industrialized and Emerging Markets. *Journal Economia Chilena* 17:48–78.

Burnside, C. 2016. Identification and inference in linear stochastic discount factor models with excess returns. *Journal of Financial Econometrics* 14:295–330.

Burnside, C. 2019. Exchange Rates, Interest Parity, and the Carry Trade. In *Oxford Research Encyclopedia of Economics and Finance*. Forthcoming: Oxford University Press.

Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo. 2011. Do Peso Problems Explain the Returns to the Carry Trade? *Review of Financial Studies* 24:853–891.

Burnside, C., M. Eichenbaum, and S. Rebelo. 2011. Carry Trade and Momentum in Currency Markets. *Annual Review of Financial Economics* 3:511–535.

Bussiere, M., M. D. Chinn, L. Ferrara, and J. Heipertz. 2018. The New Fama Puzzle. NBER Working Paper No. 24342.

Cerrato, M., N. Sarantis, and A. Saunders. 2011. An investigation of customer order flow in the foreign exchange market. *Journal of Banking & Finance* 35:1892–1906.

Cochrane, J. H. 2009. *Asset Pricing:(Revised Edition)*. Princeton, NJ: Princeton University Press.

Du, W., A. Tepper, and A. Verdelhan. 2018. Deviations from Covered Interest Rate Parity. *Journal of Finance* 73:915–957.

Engel, C. 1996. The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence. *Journal of Empirical Finance* 3:123–92.

Evans, M. D., and R. K. Lyons. 2002. Order flow and exchange rate dynamics. *Journal of Political Economy* 110:170–180.

Evans, M. D., and R. K. Lyons. 2006. Understanding Order Flow. *International Journal of Finance and Economics* 11:3–23.

Evans, M. D., and D. Rime. 2012. Micro approaches to foreign exchange determination. In J. James, I. Marsh, and L. Sarno (eds.), *Handbook of Exchange Rates*, Wiley Handbooks in Financial Engineering and Econometrics, chap. 3. Wiley.

Evans, M. D., and D. Rime. 2019. Microstructure of foreign exchange markets. Norges Bank Working Paper 6/2019.

Evans, M. D. D. 2011. *Exchange-Rate Dynamics*. Princeton, NJ: Princeton University Press.

Fama, E. F. 1984. Forward and Spot Exchange Rates. *Journal of Monetary Economics* 14:319–38.

Fama, E. F., and J. D. MacBeth. 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy* 81:607–636.

Fan, Z., J. M. Londono, and X. Xiao. 2022. Equity tail risk and currency risk premiums. *Journal of Financial Economics* 143:484–503.

Galati, G., A. Heath, and P. McGuire. 2007. Evidence of carry trade activity .

Hansen, L. P. 1982. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica* 50:1029–1054.

Hansen, L. P., and R. J. Hodrick. 1980. Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis. *Journal of Political Economy* 88:829–853.

Hansen, L. P., and R. Jagannathan. 1997. Assessing Specification Errors in Stochastic Discount Factor Models. *Journal of Finance* 52:557–590.

Harvey, C. R., and A. Siddique. 2000. Conditional skewness in asset pricing tests. *The Journal of Finance* 55:1263–1295.

King, M. R., C. L. Osler, and D. Rime. 2013. The market microstructure approach to foreign exchange: Looking back and looking forward. *Journal of International Money and Finance* 38:95–119.

Kleibergen, F., and R. Paap. 2006. Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics* 133:97–126.

Lustig, H., N. Roussanov, and A. Verdelhan. 2011. Common risk factors in currency markets. *Review of Financial Studies* 24:3731–3777.

Lustig, H., and A. Verdelhan. 2012. Exchange rates in a stochastic discount factor framework. *Handbook of Exchange Rates* pp. 391–420.

Lyons, R. K. 2001. *The microstructure approach to exchange rates*, vol. 12. Cambridge, MA: MIT Press.

Mancini-Griffoli, T., and A. Ranaldo. 2010. Limits to Arbitrage during the Crisis: Funding Liquidity Constraints and Covered Interest Parity. Swiss National Bank Working Paper 2010-14.

Meese, R. A., and K. Rogoff. 1983. Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics* 14:3–24.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf. 2012a. Carry trades and global foreign exchange volatility. *Journal of Finance* 67:681–718.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf. 2012b. Currency momentum strategies. *Journal of Financial Economics* 106:660–684.

Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf. 2016. Information flows in foreign exchange markets: Dissecting customer currency trades. *Journal of Finance* 71:601–634.

Osler, C. L. 2009. Market Microstructure, Foreign Exchange. In R. A. Myers (ed.), *Encyclopedia of Complexity and System Science*, pp. 5404–5438. Springer.

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

Sarno, L., and M. P. Taylor. 2001. The Microstructure of the Foreign-Exchange Market: A Selective Survery of the Literature. Princeton Studies in International Economics No. 89.

Shanken, J., and G. Zhou. 2007. Estimating and testing beta pricing models: Alternative methods and their performance in simulations. *Journal of Financial Economics* 84:40–86.

Villanueva, O. M. 2007. Forecasting Currency Excess Returns: Can the Forward Bias be Exploited? *Journal of Financial and Quantitative Analysis* 42:963–990.

Zhang, S. 2022. Dissecting currency momentum. *Journal of Financial Economics* 144:154–173.

Table 1: Interest-Rate and Momentum-Sorted Portfolios: Summary Statistics

Portfolio	C1	C2	C3	C4	C5	HMLC
Mean(%)	1.36	4.48	5.45	6.44	12.58	11.64
	(2.16)	(3.01)	(3.19)	(3.83)	(4.04)	(3.54)
SD	6.80	9.28	9.34	11.95	12.85	11.72
SR	0.20	0.48	0.58	0.54	0.98	0.99
Skew	-0.07	-0.88	-0.61	-1.08	-0.96	-0.78
AC1	0.07	-0.01	0.06	-0.01	-0.10*	-0.17***
Coskew1	0.42	-0.14	-0.09	-0.43	-0.38	-0.44
Coskew2	5.33	-1.50	-1.09	-7.42	-9.98	-14.59
Portfolio	M1	M2	M3	M4	M5	HMLM
Mean(%)	2.88	6.34	4.95	7.33	10.74	7.86
	(3.28)	(3.26)	(3.41)	(3.22)	(3.18)	(2.83)
SD	10.97	10.19	9.65	9.37	10.00	10.54
SR	0.26	0.62	0.51	0.78	1.07	0.75
Skew	-0.79	-0.60	-0.77	-0.66	-0.51	0.22
AC1	-0.08	-0.03	0.05	0.03	0.02	-0.16***
Coskew1	-0.34	-0.20	-0.14	-0.20	0.09	0.28
Coskew2	-7.03	-2.99	-1.72	-2.62	1.85	8.87

Note: The table reports the descriptive statistics for portfolios C1 to C5, which are sorted on the basis of short term interest rates, and portfolios M1 to M5, which are sorted on the basis of momentum (measured by previous 4-week returns). It reports statistics for the "high-minus-low" portfolios HMLC (C5–C1) and HMLM (M5–M1). It reports the annualized mean return (%) (with heteroskedasticity-consistent standard errors reported in parentheses), standard deviation (SD), Sharpe ratio (SR), skewness (Skew), and first-order autocorrelation coefficient (AC1) for each return, as well as the significance of the latter (***0.01%, **1%, *5%). We also report two measures of coskewness between the individual portfolios and the DOL portfolio. Coskew1 and Coskew2 correspond, repectively, to β_{SKS} and β_{SKD} as described in the main text.

Table 2: Exchange Rates and Order Flow for Individual Currencies

	$a_k \times 100$	$b_k \times 100$	$\bar{R}^2$		$a_k \times 100$	$b_k \times 100$	$\bar{R}^2$
AUD	-0.149	-1.148	0.055	KRW	-0.060	-1.487	0.024
	(0.086)	(0.294)			(0.069	(0.506)	
BRL	-0.122	-1.599	0.013	MXN	0.053	-1.107	0.004
	(0.098)	(0.600)			(0.065	(0.841)	
CAD	-0.066	-0.584	0.018	NOK	-0.075	-2.450	0.040
	(0.055)	(0.200)			(0.076	(0.493)	
CHF	-0.079	-0.372	0.019	NZD	-0.167	-4.234	0.099
	(0.069)	(0.112)			(0.079	(0.576)	
CZK	-0.128	-5.329	0.019	PLN	-0.098	-4.077	0.037
	(0.083)	(1.983)			(0.098)	(1.108)	
EUR	-0.174	-0.265	0.063	SEK	-0.086	0.481	0.000
	(0.064)	(0.059)			(0.076)	(0.451)	
GBP	-0.020	-0.256	0.017	SGD	-0.071	-1.039	0.038
	(0.059)	(0.085)			(0.029)	(0.257)	
HKD	0.000	-0.038	0.006	SKK	-0.141	0.979	-0.002
	(0.003)	(0.018)			(0.072)	(2.265)	
HUF	-0.039	-5.783	0.035	TRY	0.072	-4.789	0.094
	(0.096)	(1.506)			(0.085)	(0.629)	
JPY	-0.001	-0.504	0.077	ZAR	0.005	-3.833	0.070
	(0.058)	(0.087)			(0.107)	(0.568)	

Note: The table reports estimates of equation (11),

$$s_{k,t+1} - s_{k,t} = a_k + x_{k,t+1} b_k + u_{k,t+1},$$

where $s_{k,t}$ is the natural log of the exchange rate between the USD and foreign currency k, measured as foreign currency units (FCUs) per USD, $x_{k,t+1}$ is the aggregate order flow (the total value of buy orders, net of sell orders) for currency k in the interval between periods t and $t+1$, and $u_{k,t+1}$ is an error term. Heteroskedasticity consistent standard errors are reported in parentheses.

Table 3: Order-Flow Portfolios: Summary Statistics

	O1	O2	O3	O4	O5	Avg.	BMS
A) Aggregated order flow/Full sample							
Mean (%)	-7.45	1.25	9.33	10.15	19.47	5.15	27.37
	(4.13)	(3.16)	(2.96)	(3.12)	(3.68)	(3.52)	(2.66)
SD	10.94	9.73	9.51	9.16	10.30	9.34	7.31
SR	-0.68	0.13	0.98	1.11	1.89	0.55	3.74
B) Disaggregated order flow/Major currency sample							
	Asset manager						
Mean (%)	-7.36	0.74	6.42	16.36		4.04	23.72
	(3.82)	(3.61)	(3.09)	(3.19)		(3.01)	(2.62)
SD	11.25	10.79	10.18	9.72		9.11	8.61
SR	-0.65	0.07	0.63	1.68		0.44	2.75
	Hedge fund						
Mean (%)	-10.34	2.12	6.71	17.24		3.93	27.57
	(3.59)	(3.86)	(3.12)	(3.24)		(3.00)	(3.23)
SD	10.63	11.26	9.53	9.97		9.06	8.96
SR	-0.97	0.19	0.70	1.73		0.43	3.08
	Corporate						
Mean (%)	8.55	8.51	4.80	1.66		5.88	-6.89
	(3.72)	(3.19)	(3.40)	(3.15)		(3.00)	(2.42)
SD	10.40	10.35	10.40	10.26		9.01	7.25
SR	0.82	0.82	0.46	0.16		0.65	-0.95
	Private Client						
Mean (%)	22.39	11.84	-0.42	-6.31		6.88	-28.70
	(3.25)	(3.82)	(3.79)	(2.96)		(3.04)	(2.62)
SD	10.38	10.53	10.34	10.06		9.05	8.14
SR	2.16	1.13	-0.04	-0.63		0.76	-3.52

Note: For each of the portfolios O1–O5, which are sorted by contemporaneous order flow, this table reports the annualized mean excess return (with heteroskedasticity consistent standard errors reported in parentheses), standard deviation (SD) and Sharpe ratio (SR) for currencies sorted on contemporaneous order flow. Column 'Avg.' shows the average across all portfolios. Column 'BMS' (buy minus sell) reports the return of holding O5 long and O1 short. The first panel reports statistics for portfolios based on normalized aggregated order flow for the full sample of 20 currencies. The lower panels report statistics portfolios based on disaggregated order flow for a smaller sample of nine major currencies, where the disaggregation is by customer type.

Table 4: Double Sorts on Interest Rate and Order Flow: Mean Returns (%)

	Interest rate			
Order flow	Low	Medium	High	HML
Sell	-3.58	1.76	3.27	6.85
	(2.77)	(3.48)	(4.28)	(3.05)
Buy	7.79	10.38	17.75	9.96
	(2.40)	(3.10)	(4.21)	(3.90)
BMS	11.37	8.62	14.48	
	(2.09)	(1.85)	(2.72)	

Note: This table reports the annualized mean returns (with heteroskedasticity consistent standard errors in parentheses) for six double-sorted portfolios based on interest rate and the value of aggregated order flow.

Table 5: Double Sorts on Volatility Innovation and Order Flow: Mean Returns (%)

	Volatility Innovation			
Order flow	Low	Medium	High	HML
Sell	7.70	2.09	-2.26	-9.96
	(2.55)	(3.25)	(4.50)	(2.99)
Buy	14.92	10.46	5.84	-9.08
	(2.20)	(2.76)	(4.72)	(3.94)
BMS	7.21	8.37	8.09	
	(1.62)	(1.76)	(2.97)	

Note: This table reports the annualized mean returns (with heteroskedasticity consistent standard errors in parentheses) for six double sorted portfolios based on volatility innovations and the value of aggregated order flow.

Table 6: Betas of the Currency Portfolios for the DOL-HMLC Model

	α	β-DOL	β-HMLC	$\bar{R}^2$		α	β-DOL	β-HMLC	$\bar{R}^2$
C1	0.01	0.71	-0.26	0.79	M1	-0.10	0.82	0.32	0.73
	(0.02)	(0.02)	(0.02)			(0.03)	(0.04)	(0.05)	
C2	-0.01	0.94	0.00	0.82	M2	-0.01	0.92	0.17	0.82
	(0.02)	(0.02)	(0.02)			(0.03)	(0.02)	(0.04)	
C3	-0.01	0.89	0.11	0.84	M3	-0.03	0.90	0.13	0.83
	(0.02)	(0.02)	(0.02)			(0.03)	(0.03)	(0.02)	
C4	-0.05	1.04	0.29	0.84	M4	0.03	0.86	0.10	0.77
	(0.03)	(0.03)	(0.04)			(0.03)	(0.03)	(0.03)	
C5	0.00	0.71	0.74	0.94	M5	0.10	0.77	0.11	0.56
	(0.02)	(0.02)	(0.02)			(0.04)	(0.04)	(0.06)	

Note: We present estimates of the time series regressions

$$r_{it}^e = \alpha_i + z_t'\beta_i + \epsilon_{it}, \qquad t = 1, \ldots, T,$$

where r_{it}^e is the excess return of portfolio i at time t, and z_t is a vector of the two risk factors, DOL and HMLC. Estimates of α_i are scaled by 100. The portfolios are C1, C2, C3, C4 and C5 (the portfolios sorted by interest rate) as well as M1, M2, M3, M4, and M5 (the portfolios sorted by momentum of past four week returns). Standard errors are reported in parentheses. We use weekly data, from the last week of November 2001 to the fourth week of March 2012.

Table 7: Betas of the Currency Portfolios for the DOL-DVOL Model

	α	β-DOL	β-DVOL	$\bar{R}^2$		α	β-DOL	β-DVOL	$\bar{R}^2$
C1	-0.04	0.60	0.19	0.62	M1	-0.04	0.95	-0.13	0.63
	(0.02)	(0.04)	(0.07)			(0.04)	(0.07)	(0.13)	
C2	-0.01	0.94	0.10	0.83	M2	0.02	1.00	-0.04	0.78
	(0.02)	(0.02)	(0.09)			(0.03)	(0.05)	(0.11)	
C3	0.01	0.94	-0.10	0.83	M3	0.00	0.96	-0.12	0.81
	(0.02)	(0.03)	(0.08)			(0.03)	(0.03)	(0.06)	
C4	0.01	1.15	-0.25	0.78	M4	0.05	0.90	-0.10	0.76
	(0.03)	(0.06)	(0.13)			(0.03)	(0.02)	(0.06)	
C5	0.14	1.02	-0.45	0.55	M5	0.12	0.81	-0.15	0.55
	(0.05)	(0.08)	(0.20)			(0.04)	(0.05)	(0.09)	

Note: We present estimates of the time series regressions

$$r_{it}^e = \alpha_i + z_t'\beta_i + \epsilon_{it}, \qquad t = 1, \ldots, T,$$

where r_{it}^e is the excess return of portfolio i at time t, and z_t is a vector of the two risk factors, DOL and DVOL. Estimates of α_i and β-DVOL are scaled by 100. The portfolios are C1, C2, C3, C4 and C5 (the portfolios sorted by interest rate) as well as M1, M2, M3, M4, and M5 (the portfolios sorted by momentum of past four week returns). Standard errors are reported in parentheses. We use weekly data, from the last week of November 2001 to the fourth week of March 2012.

Table 8: Betas of the Currency Portfolios for the DOL-CO Model

	α	β-DOL	β-CO	$\bar{R}^2$		α	β-DOL	β-CO	$\bar{R}^2$
C1	-0.05	0.61	-0.31	0.63	M1	-0.04	0.97	0.01	0.63
	(0.02)	(0.04)	(0.06)			(0.04)	(0.07)	(0.11)	
C2	-0.01	0.94	-0.05	0.83	M2	0.03	0.99	0.20	0.79
	(0.02)	(0.03)	(0.04)			(0.03)	(0.05)	(0.07)	
C3	0.02	0.93	0.21	0.83	M3	0.00	0.96	0.12	0.81
	(0.02)	(0.03)	(0.07)			(0.03)	(0.03)	(0.07)	
C4	0.02	1.14	0.34	0.78	M4	0.06	0.89	0.19	0.76
	(0.04)	(0.06)	(0.08)			(0.03)	(0.03)	(0.08)	
C5	0.17	1.02	0.54	0.56	M5	0.14	0.80	0.28	0.55
	(0.05)	(0.09)	(0.13)			(0.04)	(0.05)	(0.10)	

Note: We present estimates of the time series regressions

$$r_{it}^e = \alpha_i + z_t'\beta_i + \epsilon_{it}, \qquad t = 1, \ldots, T,$$

where r_{it}^e is the excess return of portfolio i at time t, and z_t is a vector of the two risk factors, DOL and CO. Estimates of α_i are scaled by 100. The portfolios are C1, C2, C3, C4 and C5 (the portfolios sorted by interest rate) as well as M1, M2, M3, M4, and M5 (the portfolios sorted by momentum of past four week returns). Standard errors are reported in parentheses. We use weekly data, from the third week of January 2002 to the fourth week of March 2012.

Table 9: Betas of the Currency Portfolios for the DOL-MO Model

	α	β-DOL	β-MO	$\bar{R}^2$		α	β-DOL	β-MO	$\bar{R}^2$
C1	-0.03	0.59	0.05	0.61	M1	-0.03	0.96	-0.32	0.64
	(0.02)	(0.04)	(0.07)			(0.04)	(0.06)	(0.12)	
C2	-0.01	0.94	0.12	0.83	M2	0.02	1.00	-0.05	0.78
	(0.02)	(0.03)	(0.05)			(0.03)	(0.05)	(0.06)	
C3	0.00	0.95	0.00	0.83	M3	-0.01	0.97	0.07	0.81
	(0.02)	(0.03)	(0.05)			(0.03)	(0.03)	(0.06)	
C4	0.00	1.17	-0.05	0.77	M4	0.04	0.90	0.14	0.76
	(0.04)	(0.06)	(0.08)			(0.03)	(0.02)	(0.10)	
C5	0.14	1.06	-0.12	0.55	M5	0.12	0.82	0.32	0.56
	(0.05)	(0.08)	(0.14)			(0.04)	(0.05)	(0.11)	

Note: We present estimates of the time series regressions

$$r_{it}^e = \alpha_i + z_t'\beta_i + \epsilon_{it}, \qquad t = 1, \ldots, T,$$

where r_{it}^e is the excess return of portfolio i at time t, and z_t is a vector of the two risk factors, DOL and MO. Estimates of α_i are scaled by 100. The portfolios are M1, M2, M3, M4 and M5 (the portfolios sorted by momentum), described in the main text. Standard errors are reported in parentheses. We use weekly data, from the third week of January 2002 to the fourth week of March 2012.

Table 10: Betas of the Currency Portfolios for the DOL-CM Model

	α	β-DOL	β-CM	$\bar{R}^2$		α	β-DOL	β-CM	$\bar{R}^2$
C1	-0.03	0.60	-0.25	0.62	M1	-0.05	0.98	-0.36	0.63
	(0.02)	(0.04)	(0.10)			(0.04)	(0.07)	(0.15)	
C2	-0.01	0.94	0.09	0.83	M2	0.02	1.00	0.13	0.78
	(0.02)	(0.03)	(0.07)			(0.03)	(0.05)	(0.09)	
C3	0.01	0.94	0.21	0.83	M3	0.00	0.96	0.19	0.81
	(0.02)	(0.03)	(0.08)			(0.03)	(0.03)	(0.09)	
C4	0.01	1.16	0.28	0.77	M4	0.05	0.89	0.34	0.76
	(0.04)	(0.07)	(0.12)			(0.03)	(0.02)	(0.10)	
C5	0.15	1.04	0.39	0.55	M5	0.14	0.80	0.63	0.56
	(0.06)	(0.09)	(0.20)			(0.04)	(0.05)	(0.16)	

Note: We present estimates of the time series regressions

$$r_{it}^e = \alpha_i + z_t'\beta_i + \epsilon_{it}, \qquad t = 1, \ldots, T,$$

where r_{it}^e is the excess return of portfolio i at time t, and z_t is a vector of the two risk factors, DOL and CM. Estimates of α_i are scaled by 100. The portfolios are M1, M2, M3, M4 and M5 (the portfolios sorted by momentum), described in the main text. Standard errors are reported in parentheses. We use weekly data, from the third week of January 2002 to the fourth week of March 2012.

Table 11: Estimates of the DOL-HMLC Model

GMM Estimates				
	DOL	HMLC	R^2	HJ
b	3.85	5.93	0.39	11.73
	(4.00)	(3.78)		[0.16]
λ	0.10	0.18		
	(0.06)	(0.10)		
Fama-MacBeth Estimates				
	DOL	HMLC	R^2	χ^2_{SH}
λ	0.10	0.18	0.39	19.26
	(0.06)	(0.07)		[0.01]
KP Rank Tests				
	Stat.	d.f.	p-value	
Rank(0)	273.6	20	0.00	
Rank(1)	206.0	9	[0.00]	

Note: We present estimates of the SDF and beta representations of the DOL-HMLC model, as well as KP reduced-rank tests. The test assets are C1 to C5, the five portfolios sorted on interest rate, and M1 to M5, the five portfolios sorted on momentum. The first panel shows the estimates of the SDF coefficients, b, from first stage GMM, corresponding risk prices, λ, the cross-sectional R^2 and Hansen-Jagannathan distance (HJ). Estimates of λ are scaled by 100. The second panel shows estimates of λ obtained using the Fama-MacBeth method with no intercept. A χ^2 measure of fit is also reported. The third panel reports KP rank tests. In all panels, standard errors are reported in parentheses, and p-values in square brackets. The Shanken correction is used for the Fama-MacBeth standard errors. We use weekly data, from the last week of November 2001 to the fourth week of March 2012.

Table 12: Estimates of the Volatility (DOL-DVOL) Model

GMM Estimates				
	DOL	DVOL	R^2	HJ
b	0.58	-0.98	0.50	9.54
	(5.31)	(0.90)		[0.30]
λ	0.10	-24.61		
	(0.11)	(22.47)		
Fama-MacBeth Estimates				
	DOL	DVOL	R^2	χ^2
λ	0.10	-24.61	0.50	18.67
	(0.06)	(10.25)		[0.02]
KP Rank Tests				
	Stat.	d.f.	p-value	
Rank(0)	375.4	20	0.00	
Rank(1)	11.6	9	[0.24]	

Note: We present SDF and beta representation estimates for the DOL-DVOL model, as well as KP reduced-rank tests. The test assets are C1 to C5, the five portfolios sorted on interest rate, and M1 to M5, the five portfolios sorted on momentum. The first panel shows the estimates of the SDF coefficients, b, from first stage GMM, corresponding risk prices, λ, the cross-sectional R^2 and Hansen-Jagannathan distance (HJ). Estimates of λ are scaled by 100. The second panel shows estimates of λ obtained using the Fama-MacBeth method with no intercept. A χ^2 measure of fit is also reported. The third panel reports KP rank tests. In all panels, standard errors are reported in parentheses, and p-values in square brackets. The Shanken correction is used for the Fama-MacBeth standard errors. We use weekly data, from the last week of November 2001 to the fourth week of March 2012.

Table 13: Estimates of the DOL-CM Model

GMM Estimates				
	DOL	CM	R^2	HJ
b	-0.54	2.17	0.70	11.30
	(4.51)	(0.71)		[0.19]
λ	0.10	17.60		
	(0.06)	(5.69)		
Fama-MacBeth Estimates				
	DOL	CM	R^2	χ^2
λ	0.10	17.60	0.70	9.87
	(0.06)	(5.54)		[0.27]
KP Rank Tests				
	Stat.	d.f.	p-value	
Rank(0)	238.70	20	[0.00]	
Rank(1)	40.63	9	[0.00]	

Note: We present SDF and beta representation estimates for the DOL-CM model, as well as KP reduced-rank tests. The test assets are C1 to C5, the five portfolios sorted on interest rate, and M1 to M5, the five portfolios sorted on momentum. The first panel shows the estimates of the SDF coefficients, b, from first stage GMM, corresponding risk prices, λ, the cross-sectional R^2 and Hansen-Jagannathan distance (HJ). Estimates of λ are scaled by 100. The second panel shows estimates of λ obtained using the Fama-MacBeth method with no intercept. A χ^2 measure of fit is also reported. The third panel reports KP rank tests. In all panels, standard errors are reported in parentheses, and p-values in square brackets. The Shanken correction is used for the Fama-MacBeth standard errors. We use weekly data, from the third week of January 2002 to the fourth week of March 2012.

Table 14: Correlation Matrices of Order Flow Factors Disaggregated by Customer Segment

(a) Carry Trade Factors					
	CO	COAM	COHF	COCO	COPC
COAM	0.70	1			
COHF	0.63	0.06	1		
COCO	0.27	0.02	-0.04	1	
COPC	0.06	-0.12	-0.23	0.03	1
(b) Momentum Factors					
	MO	MOAM	MOHF	MOCO	MOPC
MOAM	0.71	1			
MOHF	0.61	0.11	1		
MOCO	0.15	-0.11	-0.16	1	
MOPC	0.08	-0.15	-0.28	0.13	1
(c) Currency Speculation Factors					
	CM	CMAM	CMHF	CMCO	CMPC
CMAM	0.70	1			
CMHF	0.61	0.10	1		
CMCO	0.22	-0.06	-0.11	1	
CMPC	0.10	-0.11	-0.29	0.14	1

Note: The table reports the correlation matrices for groups of factors defined in Sections 3 and 6.

Figure 1: Betas of Risk Factors and Average Returns

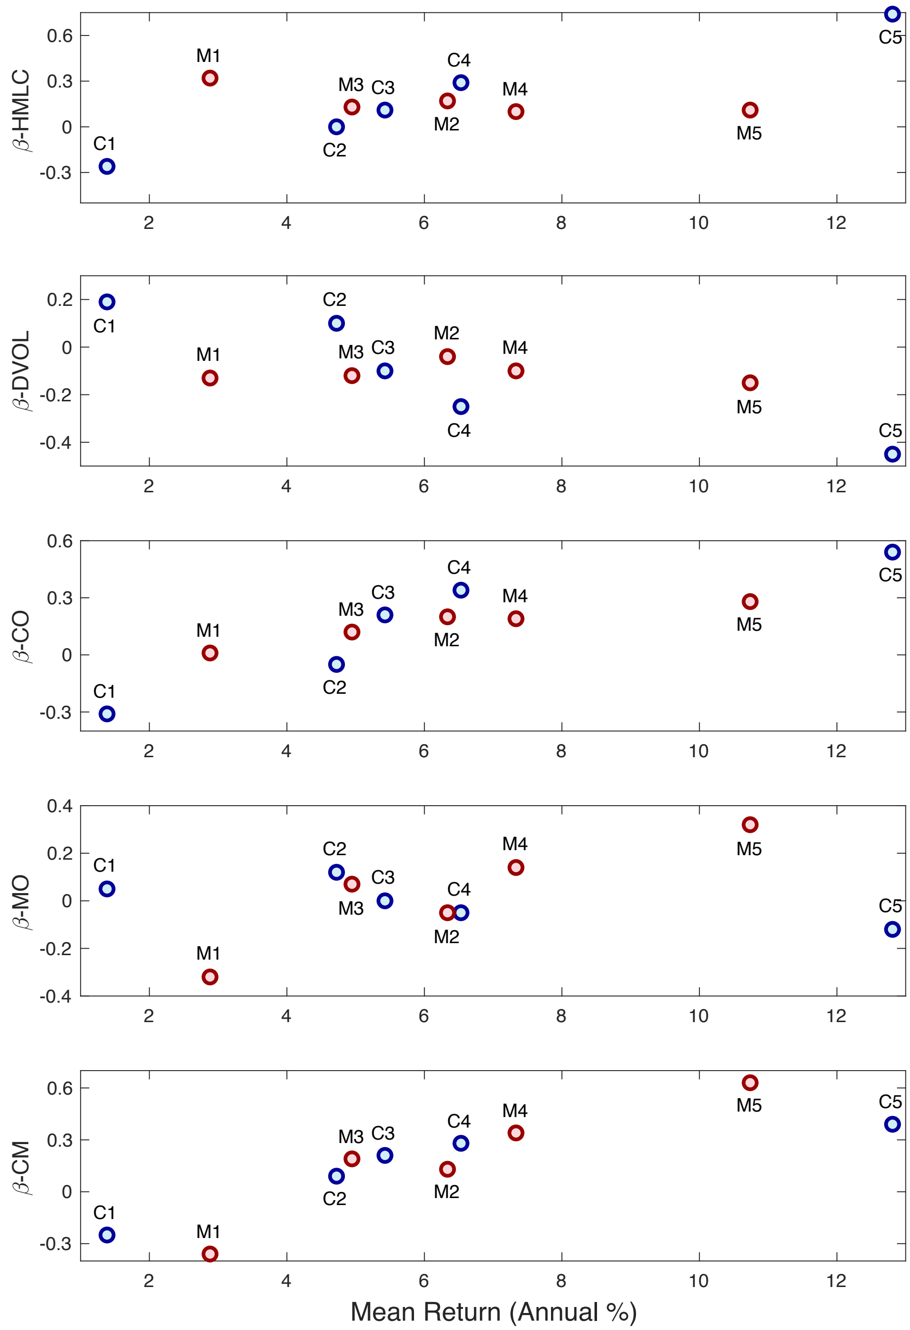

Note: This figure shows the betas of our ten currency portfolios (C1–C5 in blue, and M1–M5 in red), calculated with respect to the HMLC, DVOL, CO, MO and CM factors (when each factor is combined with DOL in a two factor model), plotted against the mean annualized excess returns of the ten portfolios.

Figure 2: Cross-Sectional Fit of the DOL-HMLC Model

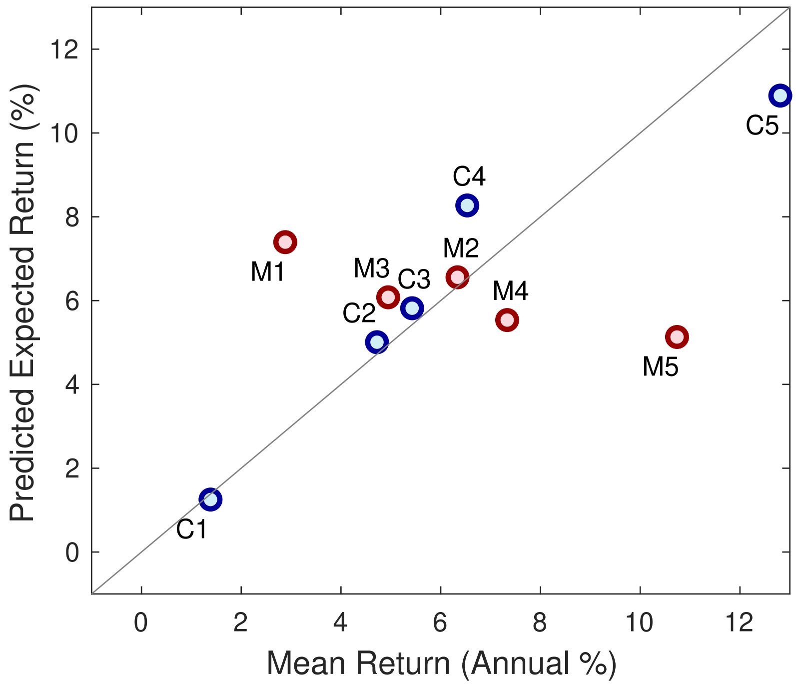

Note: This figure illustrates the cross-sectional fit of the DOL-HMLC model (see Table 11). The model-predicted expected return is plotted against the mean annualized excess returns of the ten currency portfolios (C1–C5 in blue, and M1–M5 in red).

Figure 3: Cross-Sectional Fit of the DOL-DVOL Model

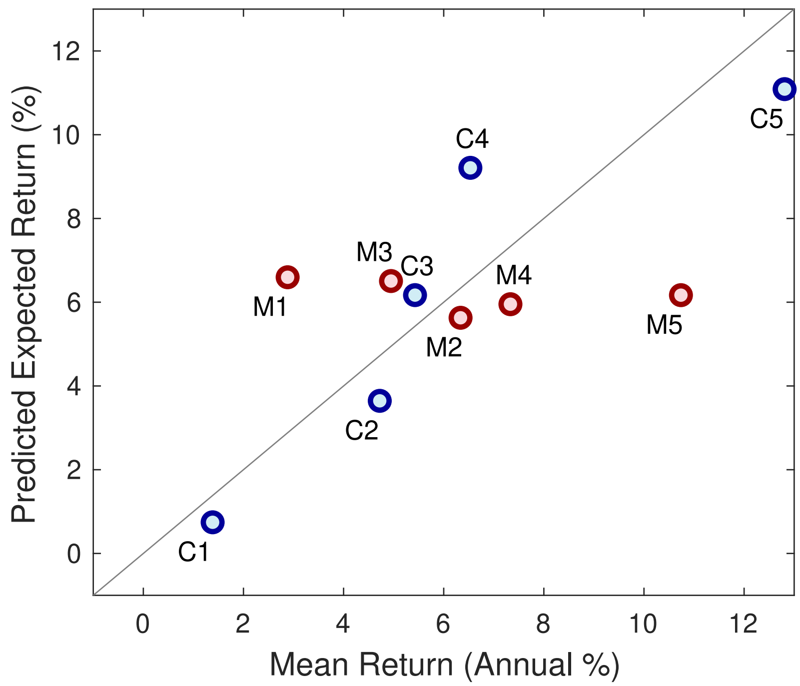

Note: This figure illustrates the cross-sectional fit of the DOL-DVOL model (see Table 12). The model-predicted expected return is plotted against the mean annualized excess returns of the ten currency portfolios (C1–C5 in blue, and M1–M5 in red).

Figure 4: Cross-Sectional Fit of the DOL-CM Model

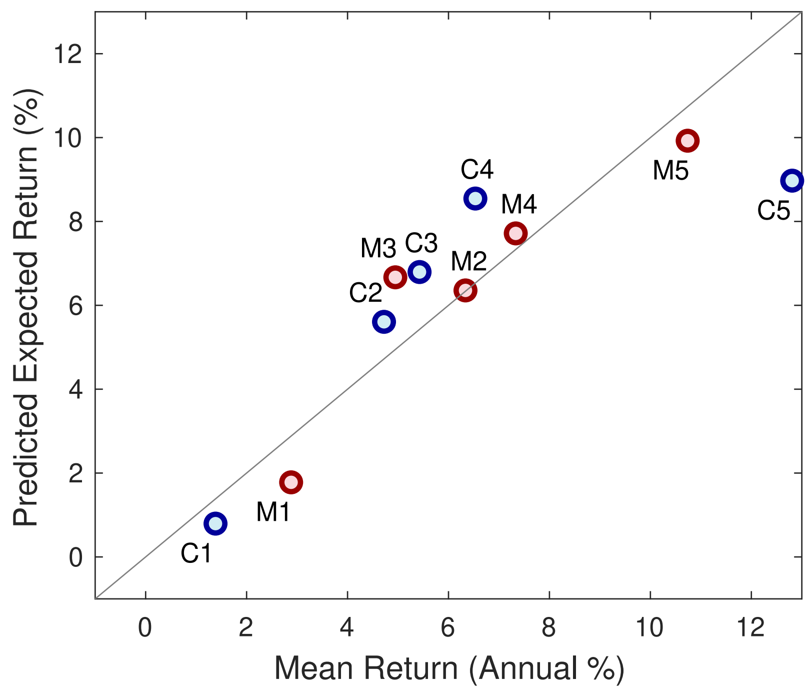

Note: This figure illustrates the cross-sectional fit of the DOL-CM model (see Table 13). The model-predicted expected return is plotted against the mean annualized excess returns of the ten currency portfolios (C1–C5 in blue, and M1–M5 in red).

Figure 5: Betas of Disaggregated Order-Flow Risk Factors

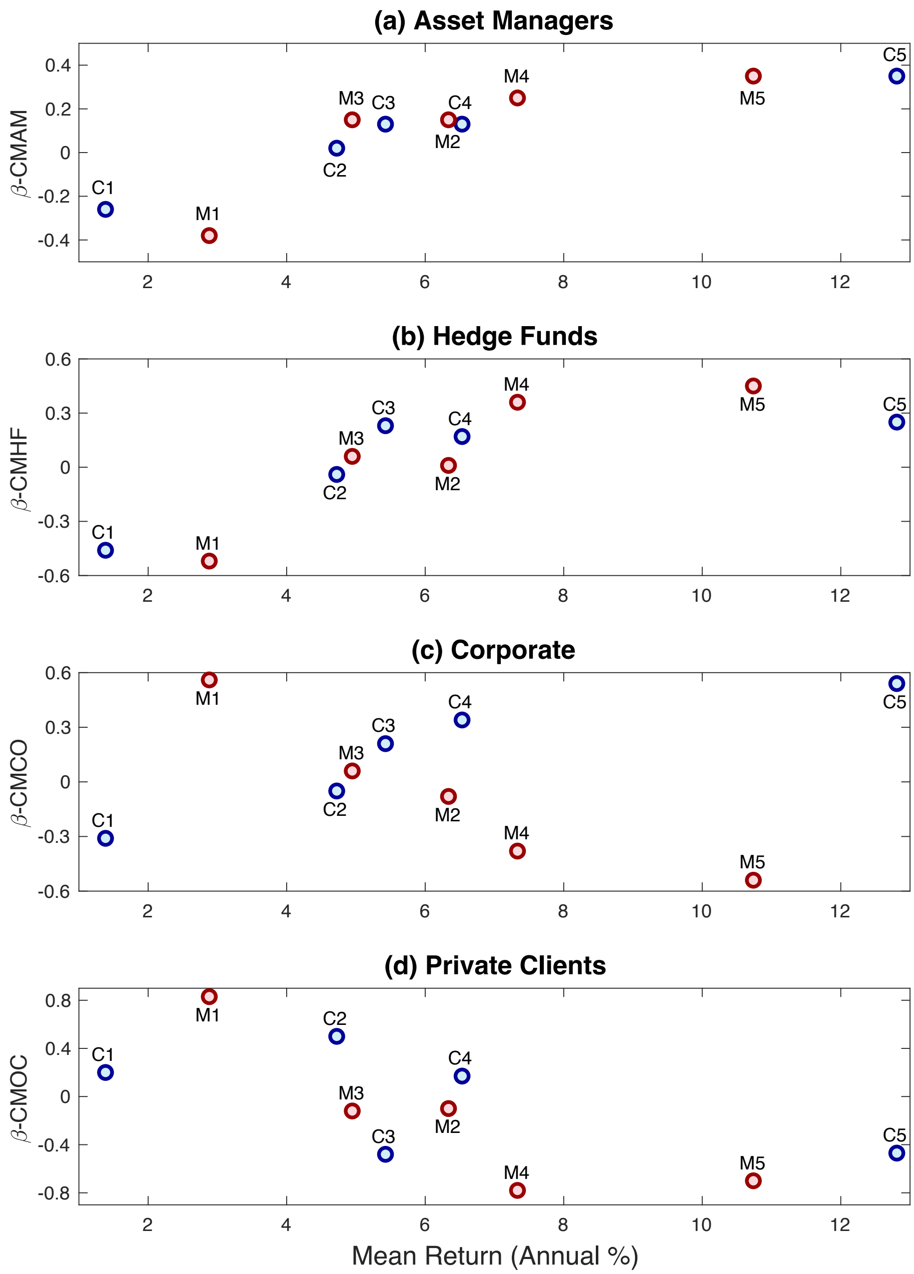

Note: This figure shows the betas of our ten currency portfolios (C1–C5 in blue, and M1–M5 in red), calculated with respect to the CMAM, CMHF, CMCO, and CMPC factors (when each factor is combined with DOL in a two factor model). The betas are plotted against the mean annualized excess returns of the ten portfolios.

Appendix

Data

Our data set consists of 20 of the most liquid currencies with the largest trading volume: the euro (EUR), Japanese yen (JPY), British pound (GBP), Swiss franc (CHF), Australian dollar (AUD), New Zealand dollar (NZD), Canadian dollar (CAD), Swedish krona (SEK), Norwegian krone (NOK), Mexican peso (MXN), Brazilian real (BRL), South African rand (ZAR), Croatian kuna (KRW), Singapore dollar (SGD), Hong Kong dollar (HKD), Turkish lira (TRY), Hungarian forint (HUF), Polish zloty (PLN), Czech krona (CZK), and Slovak koruna (SKK).

We use price quotes of spot exchange rate from the first week of November 2001 to the fourth week of March 2012. All exchange rates are quoted against US dollar, and we normalize on expressing each exchange rate as the number of FCUs per USD. The weekly and daily spot exchange rates are obtained from WM/Reuters (via Datastream). Weekly one-week forward rates are available from the same source. One-week log excess returns, defined in equation (5), are measured using the average of the bid and ask forward and spot rates.

We use a unique data set, from one of the world's largest foreign exchange dealers, that contains weekly customer order flows for the same 20 currencies from November 2001 to March 2012. We have order flow data aggregated across four types of clients—asset manager (AM), corporate clients (CO), hedge funds (HF) and private clients (PC)—for nine currencies (EUR, JPY, GBP, CHF, AUD, NZD, CAD, SEK, NOK). Asset managers and hedge funds are recognized as financial customers. Corporate and private clients are recognized as nonfinancial customers.

We believe that the order flows collected from this dealer are representative of the end-user currency demand in the foreign exchange market given that it has significant market share. The order flows measure the US dollar value of buyer-initiated minus seller initiated trades of a currency. A positive net order flow indicate net buying of foreign currency.