

Information Flows in Foreign Exchange Markets: Dissecting Customer Currency Trades

LUKAS MENKHOFF, LUCIO SARNO, MAIK SCHMELING, and ANDREAS SCHRIMPF*

ABSTRACT

We study the information in order flows in the world's largest over-the-counter market, the foreign exchange market. The analysis draws on a data set covering a broad cross-section of currencies and different customer segments of foreign exchange end-users. The results suggest that order flows are highly informative about future exchange rates and provide significant economic value. We also find that different customer groups can share risk with each other effectively through the intermediation of a large dealer, and differ markedly in their predictive ability, trading styles, and risk exposure.

JEL Classification: F31, G12, G15.

Keywords: Order Flow, Foreign Exchange Risk Premia, Heterogeneous Information, Carry Trades.

*Menkhoff is with the DIW Berlin and Humboldt University Berlin, Sarno is with Cass Business School, City University London, and the Centre for Economic Policy Research (CEPR), Schmeling is with Cass Business School, City University London, and Schrimpf is with the Bank for International Settlements. The authors would like to thank Campbell Harvey (the Editor), an anonymous Associate Editor, an anonymous referee, Alessandro Beber, Claudio Borio, Geir Bjørnes, Michael Brandt, Steve Cecchetti, Jacob Gyntelberg, Hendrik Hakenes, Joel Hasbrouck, Terrence Hendershott, Søren Hvidkjær, Gur Huberman, Alex Kostakis, Jeremy Large, Albert Menkveld, Roel Oomen, Richard Payne, Alberto Plazzi, Lasse Pedersen, Tarun Ramadorai, Jesper Rangvid, Paul Söderlind, Adrien Verdelhan, Michel van der Wel, as well as participants at several conferences, workshops, and seminars for helpful comments and suggestions. Sarno acknowledges financial support from the Economic and Social Research Council (No. RES-062-23-2340) and Menkhoff and Schmeling gratefully acknowledge financial support from the German Research Foundation (DFG). We have read the *Journal of Finance's* disclosure policy and have no conflict of interest to disclose. We agreed to allow the company that provided the data to review our paper before publication. Their feedback was helpful but did not impact our narrative. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank for International Settlements.

The foreign exchange (FX) market is the largest financial market in the world, with a daily trading volume of about five trillion U.S. dollars (Bank for International Settlements (BIS, 2013)). Also, the FX market is largely organized as an over-the-counter (OTC) market, meaning that there is no centralized exchange and that market participants can have only partial knowledge about the trades of other market participants and available liquidity in different market segments. Hence, despite its size and sophistication, the FX market is fairly opaque and decentralized because of its market structure when compared to, for example, the major equity markets. Adding to this lack of transparency, various trading platforms have been introduced and market concentration has risen dramatically over the last decade, with a handful of large dealers now controlling the lion's share of FX market turnover (see, for example, King, Osler, and Rime (2012)). In centralized, exchange-based markets, there is a single price at any point in time – the market price. In decentralized markets, by default, there is no visible common price. The FX market is the largest market of this kind.

This paper addresses several related questions that arise in this market setting. First, does customer order flow contain predictive information for future exchange rates? Answering this question is relevant for studies on market microstructure and market design, and is useful for understanding the implications of the observed shift in market concentration. Second, how does risk sharing take place in the FX market? Do customers systematically trade in opposite directions or is their trading positively correlated and unloaded onto dealers (as in, for example, Lyons (1997))? Answering these questions is also relevant for market design and provides a better understanding of the functioning of OTC markets.

Third, what characterizes different customer groups' FX trading? For example, do they speculate on trends or are they contrarian investors? And what way are they exposed to or do they hedge against market risk? Answering these questions can improve our understanding of what ultimately drives different end-users' demand for currencies and about the ecology of the world's largest financial market.

We tackle these questions empirically using a data set covering more than 10 years of daily end-user order flow for up to 15 currencies from one of the top FX dealers. The data are disaggregated into two groups of financial FX end-users (long-term demand-side investment managers and short-term demand-side investment managers) and two groups of nonfinancial FX end-users (commercial corporations and individual investors). We thus cover the trading behavior of various segments of end-users that are quite heterogeneous in their motives for market participation, informedness, and sophistication. We find that (i) order flow by end-users is highly informative about future exchange rate changes, (ii) different end-user segments actively engage in risk sharing with each other through the intermediation of a large dealer, and (iii) end-user groups show heterogeneous behavior in terms of trading styles and strategies as well as their exposures to risk and hedge factors. This heterogeneity across players is crucial for risk sharing and helps explaining the vast differences in the predictive content of flows across end-user segments that we document in this paper.

To gauge the impact of order flow on currency excess returns, we rely on a simple portfolio approach. This multi-currency framework allows for straightforward measurement of the economic value of the predictive content of order flow and is a pure out-of-sample

approach in that it only conditions on past information. Specifically, we sort currencies into portfolios to obtain a cross-section of currency excess returns, which mimics the returns to customer trading behavior and incorporates the information contained in (lagged) flows.¹ The information contained in customer trades is highly valuable from an economic perspective. We find that currencies with the highest lagged total order flows (that is, the strongest net buying pressure across all customer groups against the U.S. dollar) outperform currencies with the lowest lagged flows (that is, the strongest net selling pressure across all customer groups against the U.S. dollar) by about 10% per annum (p.a.).

For portfolios based on disaggregated customer order flow, this spread in excess returns is even more striking. A zero-cost long-short portfolio that mimics long-term demand-side investment managers' trading behavior yields an average excess return of 15% p.a., while conditioning on short-term demand-side investment managers' flows leads to a spread of about 10% p.a. Flows by commercial corporations basically generate no spread in returns, whereas individual investors' flows lead to a highly negative spread (about -14% p.a.). In sum, we find that order flow is highly informative about future exchange rates. This information is further enhanced by the non-anonymous nature of transactions in OTC markets, as trades by different categories of customers convey fundamentally different information for price movements.

What drives the predictive content in flows? We investigate three main channels. First, order flow could be related to the processing of information by market participants via the process of "price discovery." According to this view, order flow acts as the key vehicle that impounds views about (economic) fundamentals into exchange rates.² If order flow contains

private information, its effect on exchange rates is likely to be persistent. Second, there could be a price pressure (liquidity) effect due to downward-sloping demand curves (e.g., Froot and Ramadorai (2005)). If such a mechanism is at play, we are likely to observe a positive correlation between flows and prices for some limited time, followed by a subsequent reversal as prices revert to fundamental values.³ Third, we consider the possibility that order flow is linked to returns due to the different risk-sharing motives and risk exposures of market participants. For example, order flow could reflect portfolio rebalancing of investors tilting their portfolios towards currencies that command a higher risk premium. Related to this, risk-sharing could lead to the observed predictability pattern if nonfinancial customers are primarily concerned about laying off currency risk and implicitly paying an insurance premium, while financial investors are willing to take on that risk.

Discriminating between alternative explanations for the predictive content of order flow, we find clear differences across the four segments of end-users. Long-term demand-side investment managers' flows are associated with *permanent* shifts in future exchange rates, suggesting that their order flow is related to superior processing of fundamental information.⁴ In contrast, short-term demand-side investment managers' flows are associated with *transitory* exchange rate movements. This result is more in line with short-term liquidity effects than fundamental information processing. The flows of commercial corporations and individual investors seem to reflect largely uninformed trading.

Our results also point to substantial heterogeneity across customers in their trading styles and risk exposures, giving rise to different motives for risk sharing. First, we find that the trades of various end-user groups react quite differently to past returns. Long-

term demand-side investment managers tend to be “trend followers” (positive feedback traders) with regard to past currency returns. By contrast, individual investors tend to be “contrarians” (negative feedback traders). The latter finding squares well with recent findings for equity markets by Kaniel, Saar, and Titman (2008), who show that individual equity investors behave as contrarians, implicitly providing liquidity for institutional investors. Different from their results, however, individual investors do not directly benefit from serving as (implicit) counterparties of financial customers in FX markets. Second, the flows of most customer groups are negatively correlated over short to intermediate horizons, suggesting that different groups of end-users in FX markets engage in active risk sharing among each other. Thus, it is not just via the interdealer market that risk is shared in FX markets, as documented by Lyons (1997): a large dealer can provide the venue for customers to share risk due to the large size of its dealing platform, reducing the need for dealers to unload large inventories in the interdealer market. Third, we find substantial heterogeneity in the exposure to risk and hedge factors across customer segments. Long-term demand-side investment managers’ trading does not leave them exposed adversely to systematic risk, which suggests that the information in their flows is not due to risk taking but rather likely reflects superior information processing. Short-term demand-side investment managers, by contrast, are significantly exposed to systematic risk such as volatility, liquidity, and credit risk. This lends credence to the view that short-term demand-side investment managers earn positive returns in FX markets by effectively providing liquidity and selling insurance to other market participants. For nonfinancial customers there is some evidence of hedging but it is not strong enough to fully explain their negative forecast

performance arising from poor short-term market timing.

Our paper is related to prior work on the microstructure approach to exchange rates (e.g. Evans and Lyons (2002)), which suggests that order flow is crucial for understanding how information is incorporated into exchange rates. It is well known from the literature that order flow is positively associated with *contemporaneous* returns in basically all asset classes; see, for example, Hasbrouck (1991a, 1991b) for stock markets and Brandt and Kavajecz (2004) for U.S. bonds. This stylized fact also holds in FX markets, as shown by Evans and Lyons (2002) and many subsequent studies. It is less clear, however, whether order flow contains *predictive* information for exchange rates. A few papers show that FX order flow (both from interdealer and customer markets) contains information about future currency returns, but they tend to disagree on the source of this predictive power (e.g., Evans and Lyons (2005), Froot and Ramadorai (2005), Rime, Sarno, and Sojli (2010)).⁵ A few other papers fail to find robust predictive power of exchange rates by order flow in the first place, using commercially available order flow data (see, for example, Sager and Taylor (2008)). Our work is also related to a strand of recent literature that analyzes the returns to currency portfolios by investigating the predictive power of currency characteristics, such as carry or lagged returns, and the role of risk premia in currency markets.⁶

Overall, we contribute to the literature in the following ways. We are the first to show that order flow forecasts currency returns in an out-of-sample forecasting setting by directly examining currency portfolio returns based on lagged order flow. This is important as earlier papers either do not consider out-of-sample forecasting or rely on purely statistical performance measures derived from time-series forecasts of a limited number of

currency pairs (e.g., Evans and Lyons (2005), who study the DEM/USD and JPY/USD crosses). Time-series forecasts are affected by trends in exchange rates, most notably the U.S. dollar. Our portfolio procedure, by contrast, studies exchange rate predictability in dollar-neutral long-short portfolios, and it does so in an out-of-sample setting over very long time spans compared to existing FX microstructure literature. Moreover, we are the first to test whether risk exposure drives the information in customer order flows. We show how different key FX market players trade, for example, the extent to which they follow trends or behave as contrarians, and the degree to which they are exposed to systematic risk. We find strong evidence of heterogeneity in exposures and trading behavior across different groups of market participants. These findings indicate that there is significant risk sharing between financial and nonfinancial customers as well as between different groups of financial customers (long-term versus short-term demand-side investment managers) through the intermediation of a large dealer.

Taken together, these results have implications for our understanding of information flows in OTC markets. These results also add to our understanding of how risk is shared in financial markets due to different motives for trade and trading styles across end-user segments.

The rest of the paper is structured as follows. Section I describes our data, Section II presents empirical results on the predictive power of order flow, Section III empirically investigates alternative reasons for why order flow forecasts FX excess returns, and Section IV presents robustness tests. Section V concludes.

I. Data

Aggregate order flow. We employ a data set based on daily *customer* order flows for up to 15 currency pairs over the period January 2, 2001 to May 27, 2011, for a total of 2,664 trading days. Hence, in contrast to much of the earlier literature, we employ order flow from the end-user segment of the FX market and not from the interdealer market. This is important since microstructure models suggest that the information in flows stems from trading with customers and not from interdealer trading (e.g., Evans and Lyons (2002)). Order flows in our sample are measured as net buying pressure against the U.S. dollar (USD), that is, the U.S. dollar volume of buyer-initiated minus seller-initiated trades of a currency against the USD. A positive number indicates net buying pressure in the foreign currency relative to the USD. Note that order flows do not measure *trading volume* but rather net buying (or selling) pressure, as mentioned above. Aggregate order flows, that is, aggregated across customers, are available for the following 15 currencies: Australia (AUD), Brazil (BRL), Canada (CAD), the Euro (EUR), Hong Kong (HKD), Japan (JPY), Sweden (SEK), Mexico (MXN), New Zealand (NZD), Norway (NOK), Singapore (SGD), South Africa (ZAR), South Korea (KRW), Switzerland (CHF), and the United Kingdom (GBP). In the following, we refer to these flows as “total flows” since they are aggregated across all customers.

The order flows used in this paper have standard properties, similar to what has been found in other studies in this line of literature (see, for example, Froot and Ramadorai (2005)): Daily flows tend to be positively autocorrelated but the degree of autocorrelation

is very small albeit sometimes statistically significant; major currencies, such as the EUR, CHF, JPY, GBP, have much larger variation in order flows and hence a larger absolute size of order flows compared to other currencies and especially emerging markets. This is intuitive as there is much more trading in major currencies, but it also suggests that one cannot easily compare order flows across currencies and that some form of standardization is needed to make sensible comparisons.⁷ We take this into account in our empirical analysis below. Finally, aggregate order flows display high kurtosis that is largely driven by some days with extremely high (in absolute value) order flows. Eliminating these few outliers does not change our results reported below.

Disaggregated order flow. We also have access to order flows disaggregated by customer groups for the same sample period, albeit only for a subset of nine major currencies.⁸ There are four customer groups for which flows are available: long-term demand-side investment managers (LT), short-term demand-side investment managers (ST), commercial corporations (CO), and individual investors (II).⁹ Long-term demand-side investment managers comprise “real money investors,” such as mutual funds and pension funds, whereas short-term demand-side investment managers comprise other funds and proprietary trading firms. The commercial corporations’ segment includes nonfinancial corporations whereas individual investors represent trading by individuals. Hence, there is substantial heterogeneity in the motives for market participation across the four customer types, and these groups are likely to differ considerably in their degree of informedness and sophistication. Note that our data set only contains order flows based on customer-initiated trades, which means that this paper has nothing to say about the trading strategies of the dealer or any

other FX dealing bank.

Exchange rate returns and excess returns. For our empirical analysis below, we complement the above order flow data with daily spot exchange and forward rates from Reuters (available from Datastream). We denote log changes in spot exchange rates as “exchange rate returns,”

$$\Delta s_{t+1} = s_{t+1} - s_t, \quad (1)$$

where lowercase letters refer to logs and all exchange rates are quoted as the USD price of foreign currency, so that positive exchange rate returns correspond to an appreciation of the foreign currency. Hence, a positive correlation of order flows and exchange rate returns means that net buying pressure in the foreign currency (against the USD) is associated with an appreciation of the foreign currency (against the USD) and vice versa.

We also compute currency excess returns, which account for the interest rate differential in a foreign currency position. Hence, currency excess returns rx are given by

$$rx_{t+1} = s_{t+1} - s_t + (i_t^* - i_t), \quad (2)$$

where i^* denotes the foreign interest rate and i_t denotes the U.S. interest rate. Since we are working at the daily frequency in our main analysis, we need to obtain daily interest rates for all 15 countries (plus the U.S. interest rate). However, since one-day interest rates are not directly available for all countries in our sample, we employ information in forward rates to infer interest rate differentials. Interest rate differentials for horizon k are

commonly approximated by $i_{k,t}^* - i_{k,t} \approx s_t - f_{k,t}$, where $f_{k,t}$ denotes the log forward rate for horizon k of a given currency.¹⁰

II. The Value of Information in Customer Flows

A. Portfolios Conditioning on Aggregate Order Flow

We rely on a portfolio approach, mimicking the returns to customer FX trading by conditioning on lagged order flow. This provides a straightforward and intuitive assessment of how powerful order flow is in predicting currency excess returns.

As a benchmark test, we first sort currencies into portfolios based on (lagged) total order flows for each currency. Specifically, we sort currencies into five portfolios (P_1, P_2, \dots, P_5) depending on their total order flow on day t and compute portfolio excess returns (or spot exchange rate changes) for the following day. In this basic setup, portfolios are rebalanced at the end of each trading day. Note that these portfolios are computed from the viewpoint of a U.S. investor as each individual portfolio consists of a short position in USD and a long position in a basket of foreign currencies. Taking the return difference between any two portfolios $P_j - P_i$ thus gives the return of a portfolio short in the basket of foreign currencies in P_i and long in the basket of currencies in P_j , so that the USD component cancels out and the long-short portfolio is dollar-neutral by construction.

Standardizing order flows. Before sorting currencies into portfolios, we need to make sure that order flows are comparable across currencies. As the absolute size of order flows

differs across currencies it is not sensible to sort currencies based on raw order flows. To allow for meaningful cross-currency comparisons, it is necessary to standardize flows. We do this by dividing flows by their standard deviation to remove the difference in absolute order flow sizes across currencies,

$$\tilde{x}_{j,t}^R = \frac{x_{j,t}}{\sigma(x_{j,t-59;t})}, \quad (3)$$

where $\tilde{x}_{j,t}^R$ denotes order flow standardized over a rolling window and $x_{j,t}$ denotes the raw order flow. In our baseline results, we compute the standard deviation of flows via a rolling scheme over a 60-day window. Robustness tests based on alternative approaches to standardize flows are reported in a separate Internet Appendix.¹¹

Portfolio excess returns. Table I shows average annualized excess returns for order flow portfolios (P_1, P_2, \dots, P_5), where P_1 contains the three currencies with the lowest lagged standardized order flow and P_5 contains the three currencies with the highest lagged standardized order flow. Hence, P_5 can be thought of as a portfolio of currencies with the highest buying pressure, whereas P_1 refers to a portfolio with the strongest selling pressure. Column “Av.” shows average returns across all currencies in the cross-section and column “BMS” denotes a portfolio that is long in P_5 and short in P_1 (“Buying Minus Selling” pressure). We report returns for the full sample period from January 2001 to May 2011.¹²

To get started, Panel A of Table I reports results for the sample of all 15 markets (T15) as well as the subsample of nine developed markets (T9); for the T9 subsample we form only four portfolios rather than five to ensure we always have two currencies in the corner

portfolios. We observe a strong increase in average excess returns as we move from the portfolio of currencies with low buying pressure, P_1 , to the one with high buying pressure, P_5 (or P_4 for the T9 sample). The spread in excess returns between the high buying pressure portfolio and the low buying pressure portfolio, that is, the excess return of the BMS portfolio, is economically large (10.31% and 12.43% p.a., respectively) and statistically highly significant. Similarly, the Sharpe Ratios (p.a.) of the two BMS portfolios of 1.26 and 1.45 are large and also point toward high economic significance. Thus, order flows carry significant information for future currency excess returns, as captured by our dollar-neutral out-of-sample trading strategy that only conditions on real-time information.

TABLE I ABOUT HERE

Table IA.II in the Internet Appendix shows results for the other standardization schemes and by subsample. We find that our results are equally strong across various subperiods. Table IA.III in the Internet Appendix repeats this exercise for exchange rate changes instead of excess returns. The results in that table clearly show that the patterns in average spot exchange rate changes across portfolios are at least as strong as those for average excess returns. Hence, order flow is informative about future spot rates and not about interest rate differentials. Figure IA.1 in the Internet Appendix plots the cumulative excess returns to the BMS portfolios.

Tests for return monotonicity. The columns “MR” and “Up” in Table I report tests for return monotonicity (Patton and Timmermann (2010)), that is, whether there is a significantly increasing or decreasing pattern of average excess returns when moving from

the portfolio of low buying pressure (P_1) to the one with high buying pressure (P_5).¹³ These tests go beyond the standard t -test of a zero BMS portfolio return since they take into account the entire cross-sectional pattern. This is interesting since one would intuitively expect an increasing pattern of average portfolio excess returns when moving from P_1 to P_5 if order flow is truly informative about future excess returns. This prediction is significantly borne out in the data for both the “MR” and the “Up” tests. Hence, there is strong evidence for a significant relationship between order flow and future excess returns.

B. Portfolios Conditioning on Disaggregated Order Flow

If superior information processing or genuine forecasting ability drive our results above, one would expect clear differences in the forecasting power of different customers’ order flows, depending on the groups’ characteristics (see, for example, Evans and Lyons (2007)). Specifically, one would expect to see superior information processing in flows of financial customers, given that nonfinancial players do not specialize in FX trading as their core activity. To investigate this, we now build portfolios based on our disaggregated data for customer flows. We closely follow the earlier approach with the exception that we build only four portfolios (rather than five) here since we have disaggregated flows for only nine currencies and want to have a minimum of two currencies per portfolio.

Table I, Panel B reports results for the four customer groups (long-term demand-side investment managers (LT), short-term demand-side investment managers (ST), commercial corporations (CO), and individual investors (II)). The results are clear-cut. Long-term demand-side investment managers’ net buying or selling pressure for currencies is the most

informative about subsequent exchange rate behavior. Conditioning on long-term demand-side investment managers' flows generates a cross-sectional spread in excess returns of 15% p.a., followed by short-term demand-side investment managers with a spread of about 10%. In stark contrast, the flows of commercial corporations and individual investors actually generate a *negative* spread in portfolio excess returns of about -4% and -14% , respectively.¹⁴ The results point towards substantial differences in customers' predictive information. The latter is underscored by the large spread in (annualized) Sharpe Ratios of BMS portfolios across customer groups. Long-term demand-side investment managers' BMS portfolio yields a Sharpe Ratio of 1.79, whereas individual investors' BMS portfolio has a Sharpe Ratio of -1.55.¹⁵

As above, we also present p -values for tests of return monotonicity. Since the order flow of corporations and individual customers negatively forecasts returns, in these cases we modify the MR test to check for a monotonically decreasing pattern. Results from these tests corroborate the simple t -tests for the BMS portfolios. There is a monotonically increasing pattern in average excess returns for portfolios based on long-term demand-side investment managers' and short-term demand-side investment managers' flows that is highly significant. By contrast, we find a monotonically decreasing pattern in average excess returns for portfolios based on individual investors' flows, and marginally significant evidence for a decreasing pattern in portfolios based on commercial corporations' flows.

Taken together, the results show that not all order flow is equal in terms of its information content for exchange rates. Instead, financial customers' flows (long-term demand-side investment managers and short-term demand-side investment managers) account for the

positive relation between lagged flows and future exchange rate returns uncovered in the previous section. Flows of commercial corporations are more or less uninformative, and individual investors' flows even forecast returns in the wrong direction. The latter finding of poor trading performance and market-timing skills by individual investors is in line with earlier evidence for stock markets that shows individuals tend to lose money from trading (e.g., Grinblatt and Keloharju (2000), Hvidkjaer (2008), Barber et al. (2009)). Using total end-user order flow masks these differences and might even lead to incorrect inferences about the link between flows and returns. In a nutshell, what matters for the relation between end-user order flows and future returns is disaggregated data, since the information content of flows for future returns varies markedly across customer groups.

The middle and lower panels of Figure IA.1 plot cumulative returns for all four customer groups. It can directly be seen that returns are very different across customer groups, even when comparing, for example, long-term demand-side investment managers and short-term demand-side investment managers. Both groups' BMS portfolios generate significant excess returns, but returns for short-term demand-side investment managers are much more volatile than those of long-term demand-side investment managers. We investigate possible sources of these different return behaviors below.

C. Marginal Predictive Content of Flows at Longer Horizons

Our analysis so far focuses on the relation between order flows and returns over the subsequent trading day. An interesting question that arises, however, is whether the information contained in order flow quickly decays or is useful for forecasting returns over more than

one trading day.

To examine the marginal predictive content of flows, we form portfolios as in the analysis above but we now allow for a longer lag between the order flow signal and portfolio formation. Table II presents the results for lags of $0, 1, 2, \dots, 9$ days. To be more specific, a lag of zero days means that flows of trading day t are used to predict returns of day $t + 1$ (thus reproducing the BMS returns from Table I above), whereas a lag of, say, two days means that flows of day t are used to forecast returns of trading day $t + 3$.

TABLE II ABOUT HERE

The results in Table II show that order flow appears to be most informative for the first two to three days after portfolio formation, with the information in flows becoming insignificant afterwards. Hence, the information contained in daily flows is fairly short-lived and is impounded into exchange rates relatively quickly. This finding is in contrast to, for example, Evans and Lyons (2005), who study a shorter and smaller sample and find that times-series predictability of returns by order flow increases at longer horizons when judged from statistical metrics of forecast evaluation. This contrast in results also highlights the importance of not assessing the predictive power of order flow based only on purely statistical measures, as statistical evidence of exchange rate predictability in and of itself does not guarantee that an investor can earn profits from a trading strategy that exploits this predictability.

III. What Drives the Predictive Power of Flows?

A. *Permanent vs. Transitory Forecast Power of Flows*

To better understand the forces driving our results above, we next investigate whether order flow forecasts returns because it signals permanent shifts in spot exchange rates or whether it merely forecasts temporary movements that are eventually reversed after some time. The question of whether order flow has a permanent or transitory effect on prices is a central one in earlier microstructure literature (see Hasbrouck (1991a, 1991b)). A transitory movement would be interpreted as suggesting that order flow effects are merely due to short-term liquidity or price pressure effects that eventually die out, whereas a permanent movement in spot rates would indicate that order flow conveys information about fundamentals.¹⁶ More specifically, a permanent price impact would indicate that order flow is related to *changes in expectations about fundamentals* given the daily frequency we are working with. Since we find substantial heterogeneity with regard to the forecasting power of different customer groups' order flows, the question of whether all (or some) customers' flows signal information relevant for permanent changes in FX rates or whether some customer groups' flows simply exert price pressure and liquidity effects is of interest.

To this end, we employ our portfolio sorts framework as above but now track *cumulative* exchange rate returns to BMS portfolios for overlapping periods of 30 trading days after portfolio formation. This approach yields a direct estimate of how spot rates move after experiencing intensive buying or selling pressure from customers.

Figure 1 illustrates the persistence of the predictive content of order flow. The solid lines show the cumulative excess returns (in basis points), whereas the shaded areas show 95% confidence intervals based on a moving-block bootstrap with 1,000 repetitions. Total flows for all 15 currencies (T15) forecast a permanent change in spot rates that is statistically significantly different from zero. Exchange rates with the highest net buying (selling) pressure appreciate (depreciate) against the USD for approximately three days. Currency returns on the BMS portfolios increase by about 15 basis points over this period, and afterwards the effect of the order flow signal levels out. Importantly, these findings suggest that order flow conveys information and its impact on exchange rates is not reversed.

FIGURE 1 ABOUT HERE

This picture changes when looking only at the nine developed currencies. Here, we observe the same increasing pattern initially, followed by a subsequent partial reversal. After approximately 25 to 30 trading days, about one-half of the initial impact of 15 basis points is reversed and the confidence interval includes zero. Hence, there is much less evidence that order flow conveys information about fundamentals when only looking at major developed markets. This finding makes sense, however, since the major currency markets are most probably more researched and more efficient than smaller currency markets, so that the scope for superior information processing is reduced.¹⁷

As a natural next step, we reexamine this question for disaggregated order flows (lower panels of Figure 1). The results are clear-cut. The only end-user group with a statistically significant permanent price impact is that of long-term demand-side investment managers.

Short-term demand-side investment managers' trading has a positive but transitory impact, commercial corporations have no impact at all, and individual investors have a transitory negative impact. Given our finding for the total flows of the nine major currencies above, it is interesting to see that long-term demand-side investment managers' flows are indeed associated with permanent spot rate changes. This suggests that the order flow of long-term demand-side investment managers is likely related to the processing of fundamental information whereas that of short-term demand-side investment managers corresponds to short-lived information that is less strongly related to fundamental information. Similarly, it is reasonable that the negative relation between individual investors' flows and future spot rates dies out over time.

These findings suggest that the order flows of different end-user groups embed different information for future exchange rates. These differences can arise either because they are based on different mechanisms to process information or because of different trading motives and hedging needs. To explore these possibilities further, and thus shed light on the observed differences in end-user order flows, we investigate the drivers of order flow in more detail below.

B. Risk-Sharing among Foreign Exchange End-Users

The analysis above suggests that long-term demand-side investment managers' order flows are related to the processing of fundamental information that is quickly and permanently impounded into prices, whereas the other customer groups' order flows are not. A potential explanation is that risk sharing among market participants drives our results, at least in

part. A risk-sharing story implies that we observe customers systematically trading in opposite directions and that their portfolios load on different sources of systematic risk. We investigate these issues below.

Portfolio returns in event time. We first provide a more detailed look at the return behavior around portfolio formation dates to better understand differences in customer groups. Figure 2 shows the average annualized BMS excess return for the five days prior to portfolio formation (days $-5, -4, \dots, -1$), the day of portfolio formation (day 0), and the first ten days after portfolio formation (days $1, 2, \dots, 10$). Shaded areas correspond to 95% confidence intervals based on Newey and West (1987) standard errors. Note that these returns, unlike in Figure 1, are not cumulative.

FIGURE 2 ABOUT HERE

Two results stand out. First, long-term demand-side investment managers tend to be trend followers in that they exert buying (selling) pressure in currencies that recently appreciated (depreciated). Conversely, individual investors tend to trade against the trend, that is, they react upon past returns in a contrarian fashion. The pattern for short-term demand-side investment managers and commercial corporations is less clear. Second, formation-day returns (day 0) are significantly different from zero for all four customer groups. However, short-term demand-side investment managers (positive) and individual investors (negative) have the largest contemporaneous returns in absolute value, indicating that either their trades heavily drive exchange rates or their trades are heavily triggered by returns (e.g., via stop-loss and stop-buy orders).

Overall, these findings suggest that customer groups' trading positions at least partly offset each other, as long-term demand-side investment managers and individual investors clearly differ in terms of their trend-following behavior. This finding is different from equity markets, for which Kaniel, Saar, and Titman (2008) find that individual investors tend to be contrarian traders but experience subsequent positive returns, presumably due to implicitly providing liquidity to institutional investors. In our data, we find similar contrarian behavior of individual investors, but this trading behavior does not yield positive returns on average.

Flow correlations over longer horizons. Given these findings, we next look at the correlations among customer groups' flows directly. While there is little contemporaneous correlation in flows, as we note above (see Table IA.I in the Internet Appendix), it is nevertheless interesting to look at flows over longer horizons to find out if customer groups tend to trade in the same or opposite directions. For a risk-sharing explanation to make sense, we would expect to see negative flow correlations between customer groups at some horizons.

Figure 3 plots contemporaneous correlations between standardized flows of different customer groups for horizons of one to 60 days (using overlapping observations), where the shaded areas correspond to 95% bootstrap confidence intervals. The two financial customer groups (long-term and short-term demand-side investment managers) tend to trade in opposite directions over very short horizons but in the same direction over the longer run. Moreover, all correlations between financial and nonfinancial customers are significantly negative at all horizons, while there is no significant correlation between flows of

the nonfinancial customer groups. These results are generally in line with a risk-sharing story whereby financial players trade in the opposite direction of nonfinancial market participants. This finding is interesting because the perception in the literature is that risk sharing only takes place in the interdealer market (see, for example, Lyons (1997)) where dealers quickly lay off their accumulated inventory from customer orders. Our results indicate that risk sharing can also take place in the customer market due to the negative correlation of the order flows of different market segments.

FIGURE 3 ABOUT HERE

Drivers of flows. As a natural next step we seek to provide a better understanding of the drivers of end-user order flows and shed light on the source of the negative flow correlations discussed above. First, we examine whether the flows of some customer groups systematically lead the flows of other groups. Second, we study whether customers' flows differ in their response to lagged asset returns in other key asset classes. In this context we are interested in the possible effects of portfolio rebalancing on the end-user demand for currencies (Hau and Rey (2004)). To investigate this question, we run panel regressions of order flows on lagged flows and further explanatory variables, such as interest rate differentials ($i_t^* - i_t$), lagged exchange rate changes over one and 20 days ($\Delta s_t, \Delta s_{t-1;t-20}$),

lagged stock returns $(r_t^{eq}, r_{t-1;t-20}^{eq})$, and lagged bond returns $(r_t^b, r_{t-1;t-20}^b)$,

$$\begin{aligned}
OF_{j,t+1}^c = & \alpha + \beta_{LT}OF_{j,t}^{LT} + \beta_{ST}OF_{j,t}^{ST} + \beta_{CO}OF_{j,t}^{CO} + \beta_{II}OF_{j,t}^{II} \\
& + \gamma_1(i_{j,t}^* - i_t) + \gamma_2\Delta s_{j,t} + \gamma_3\Delta s_{j,t-1;t-20} \\
& + \gamma_4r_{j,t}^{eq} + \gamma_5r_{j,t-1;t-20}^{eq} + \gamma_6r_{j,t}^b + \gamma_7r_{j,t-1;t-20}^b + \varepsilon_{j,t+1},
\end{aligned} \tag{4}$$

where c denotes one of the four customer groups, j denotes currencies/countries, and $\varepsilon_{j,t+1} = e_{t+1} + u_j + \epsilon_{j,t+1}$ includes both cross-sectional and time fixed effects. Standard errors are clustered by currency pair. We use benchmark 10-year government bonds and country equity indices from Datastream for bond and stock returns. The frequency is daily.

Results from these regressions are reported in Table III. For each customer group we report one specification that only includes lagged flows and one that additionally includes interest rate differentials and lagged returns.¹⁸ Looking first at the specifications that only include lagged flows, we find that the flows of long-term demand-side investment managers are significantly related to the flows of the other groups. These results (akin to simple Granger causality tests) corroborate the notion that long-term demand-side investment managers tend to trade very differently from, and indeed in the opposite direction of, non-financial customers. Flows of short-term demand-side investment managers do not load significantly on lagged flows of any group, again indicating that long-term demand-side investment managers and short-term demand-side investment managers behave quite differently. The flows of commercial corporations are positively driven by own lagged flows and lagged flows of individual investors, whereas flows of individual investors are signifi-

cantly negatively related to lagged short-term demand-side investment managers' flows and are significantly positively autocorrelated. In sum, there are numerous interrelationships between customer flows and their lags, but it may be overambitious to interpret them in any structural way.

TABLE III ABOUT HERE

When we include lagged returns as additional regressors, we find that long-term demand-side investment managers trade against the interest rate differential, whereas commercial corporations trade with the interest rate differential. Surprisingly, flows of short-term demand-side investment managers (and individual investors) are not affected by the interest differential, suggesting that on average carry trading does not drive their flows in our sample. Results for lagged exchange rates indicate that long-term demand-side investment managers are trend followers (positive feedback traders), whereas individual investors can be described as contrarians (negative feedback traders). Long-term demand-side investment managers' flows also react significantly positively to lagged equity returns, whereas individual investors' flows are positively driven by lagged bond returns. Hence, investors tend to increase their position in a currency (against the USD) when the country's stock market return has been high (long-term demand-side investment managers) or when government bond prices have been increasing (individual investors). These results do not suggest that order flows are driven by portfolio rebalancing in the sense that investors sell a currency in response to rising equity or bond prices in the country (see, for example, the mechanism described in Hau and Rey (2004)). However, the results strongly support the

notion that flows of different groups are driven in part by the returns of other asset classes. The results also show that the factors that drive flows clearly differ across end-user groups.

C. Differences in Risk Exposures

Finally, we investigate whether differences in risk exposures can account for BMS return patterns across FX end-users. A risk channel could explain the observed BMS excess returns if long-term demand-side investment managers and short-term demand-side investment managers tilt their portfolios towards risky currencies and earn a risk premium whereas commercial corporations and individual investors tilt their portfolios towards safe currencies and earn low or even negative returns.

Since there are many possible sources of systematic risk in our case, we consider an augmented version of the Fung and Hsieh (2001, 2002, 2004) multi-factor model as the basis for these risk adjustments. The Fung-Hsieh model has served as the workhorse for understanding risk exposures in the literature (see, for example, Patton and Ramadorai (2013)). The model relies on various U.S. equity market and bond market factors and also includes the returns on trend-following strategies to capture exposures to nonlinear option-like payoffs that are quite typical of hedge funds. The trend-following factors are constructed from portfolios of lookback straddles in various asset classes. We modify the framework to make it amenable to an analysis focused on the FX market and to allow for conditional exposures (e.g., Ferson and Schadt (1996), Patton and Ramadorai (2013)). The regression that serves as the basis of these tests takes the form

$$rx_{p;t} = \alpha + \sum_{k=1}^K \beta_k F_{k;t} + \sum_{j=1}^J \theta_j r_{m;t} \cdot z_{j;t-1} + \epsilon_t. \quad (5)$$

The set of factors F_t includes the excess return on the U.S. equity market (r_m), the change in the yield spread of U.S. long-term bonds (ΔTS), and changes in credit spreads (ΔDF). It further includes returns on portfolios of lookback straddles for FX futures and interest rate futures, denoted by $PTFS_{FX}$ and $PTFS_{IR}$, respectively. We augment this subset of factors from Fung and Hsieh (2004) with additional factors intended to capture FX-related risk. We include the dollar risk factor (DOL) and the carry factor (HML_{FX}) of Lustig, Roussanov, and Verdelhan (2011) as well as a factor-mimicking portfolio of global FX volatility (VOL_{FX}) from Menkhoff et al. (2012a). Following Patton and Ramadorai (2013), we also allow for conditional risk exposures by interacting the equity market factor, $r_{m;t}$, with lagged conditioning variables, $z_{j;t-1}$. In particular, we consider (a) changes in the TED spread (Brunnermeier, Nagel, and Pedersen, (2009)), (b) changes in the VIX (Whaley (2000)), and (c) the change in the three-month T-bill rate.

To keep the analysis tractable and to avoid overfitting, we perform model selection of the space of risk factors. Ideally, we want to explore the same set of factors for each of the customer segments to be able to compare the exposures across customers and learn about differences that can explain the variation in BMS excess returns. However, as financial and nonfinancial customers are likely to be very different, we focus on long-term demand-side investment managers versus short-term demand-side investment managers in the first set of results and individual investors versus commercial corporations in the second set

of results. More specifically, we perform model selection over a two-equation seemingly unrelated regression (SUR) for long-term demand-side investment managers' and short-term demand-side investment managers' BMS returns, and a separate model selection for a SUR for commercial corporations and individual investors.

The results are in Table IV. Panel A shows results for linear models, whereas Panel B allows for conditional market exposures. We report the four best-performing models with a maximum of three factors included in the regression. The best linear model in Panel A picks global FX volatility (VOL_{FX}) as the single factor. Other model specifications that also perform well tend to incorporate the trend-following factors as well as term spread and default spread changes. Interestingly, when comparing long-term demand-side investment managers' and short-term demand-side investment managers' exposures to these factors, we find that the signs are always opposite. While long-term demand-side investment managers' BMS returns load positively on FX volatility shocks, trend-following factors, and changes in the default spread, short-term demand-side investment managers load negatively on these factors. This means that long-term demand-side investment managers' FX trading positions tend to perform well in periods of market-wide stress and when there are large returns to following trends (which happens to be in volatile periods, when markets trend more). Short-term demand-side investment managers' FX trading positions, however, are adversely exposed to systematic risk and market distress. These results are quite striking as they indicate that long-term demand-side investment managers have very different FX trading behavior and exposure to systematic risk than short-term demand-side investment managers.¹⁹

TABLE IV ABOUT HERE

Allowing for conditional exposures by adding interaction terms between market returns (r_m) and lagged changes in TED spreads and the VIX (Table IV, Panel B) leaves the main factors chosen largely unchanged but tends to improve the model fit. The results reported in Panel B thus corroborate previous results that trading from long-term and short-term demand-side investment managers is very different and that their FX trading positions are differently exposed to market stress.²⁰

TABLE V ABOUT HERE

Since some of our risk factors in the above regressions are not returns (e.g., changes in yield spread and the default spread), the intercepts cannot be interpreted as a risk-adjusted return. We therefore re-run this analysis after replacing the non-return factors by their factor-mimicking portfolios. The results, reported in Table IA.XXIII in the Internet Appendix, paint a very similar picture. We find that the intercepts for financial customers are large and significant, ranging from 0.71% to 1.46% p.m. We also add two further FX-specific variables to the menu of potential conditioning variables: changes in the average forward discount (AFD) across countries (Lustig, Roussanov, and Verdelhan (2014)), which captures interest rate differentials, and changes in global FX volatility (Menkhoff et al. (2012a)). The results, provided in Table IA.XXIV of the Internet Appendix, indicate that incorporating FX volatility tends to drive out the VIX as a conditioning variable, whereas the AFD does not appear in the top model specifications. Thus, from an economic perspective, our main results are largely unchanged when considering these two FX

conditioning variables.

We repeat the analysis above for nonfinancial customers' BMS portfolios as well. The results are shown in Table V. As might be expected, we find that risk exposures do not matter as much for nonfinancial customers. Nonetheless, we do find evidence of a negative equity market exposure for both groups (Panel A), which increases (decreases) following increases in the TED spread for individuals (commercial corporations). Moreover, there is some evidence that the individual investors' BMS portfolio has positive exposure to changes in credit spreads.

IV. Additional Tests and Robustness

We provide extensive robustness checks for all our main results. These tests are briefly described below. More detailed results are reported in the Internet Appendix.

Transaction costs. An interesting question is whether the BMS returns remain large after accounting for transaction costs. To examine this question, we compute *net* excess returns for BMS portfolios by adjusting for bid-ask spreads.²¹ We investigate returns to strategies with varying portfolio rebalancing frequencies to balance the effects of transaction costs and using the most recent information. Figure IA.3 in the Internet Appendix presents the results for rebalancing frequencies from one to 10 days. The dashed lines give average excess returns (p.a.) and 95% confidence intervals for excess returns before transaction costs to show the effect of different rebalancing periods. The solid line and shaded area give average net excess returns (p.a.) and 95% confidence intervals when taking transaction

costs into account. We find that average excess returns are significantly different from zero for all rebalancing horizons and economically attractive even for short frequencies.

Panel regressions. We run panel regressions of currency returns on order flow to control for other possible determinants of currency excess returns as well as cross-sectional and time fixed effects. Specifically, we run panel regressions of the form

$$rx_{j,t+1} = \beta_c OF_t^c + \gamma_1(i_{j,t}^* - i_t) + \gamma_2 rx_{j,t} + \gamma_3 rx_{j,t-60;t-1} + \varepsilon_{j,t+1}, \quad (6)$$

where j ($1, \dots, N$) indexes currencies, rx denotes currency excess returns, OF^c denotes the order flow of customer group c , $(i_{j,t}^* - i_t)$ denotes interest rate differentials (carry), and rx_t and $rx_{t-60;t-1}$ denote lagged excess returns over the prior trading day and the average over the past 60 trading days, respectively.²² The error term is given by $\varepsilon_{j,t+1} = e_{t+1} + u_j + \epsilon_{j,t+1}$ and thus captures time and cross-sectional fixed effects (we also report results without fixed effects below). Standard errors are clustered by currency pair. These panel regressions are based on individual currency returns and not on portfolio returns.

The results reported in Table IA.VI corroborate our findings based on the portfolio approach above. In particular, the results show that order flows of financials positively predict future excess returns, whereas flows of nonfinancial end-users negatively forecast returns. Importantly, the predictive relation between lagged order flow and future FX excess returns remains very strong even when controlling for two common predictors of returns in FX markets, namely, interest rate differentials (carry trade) and (short-term) currency momentum (Menkhoff et al. (2012b)).

Standardizing flows. We next check whether our results are robust to other sensible choices of standardizing flows. First, in Tables IA.VII and IA.VIII in the Internet Appendix, we check whether standardizing flows over longer horizons of one and three years produces similar results. They do. Second, we measure flows relative to total currency trading volume (obtained from the BIS FX triennial surveys).²³ The results reported in Table IA.IX also indicate significant predictability of returns by order flows. Third, we standardize flows by additionally demeaning flows over the rolling window (Table IA.X) and we form portfolios that take positions in all available currencies with weights determined by the magnitude of order flow (Table IA.XI). Our results remain robust.

Longer horizons. We also check whether order flows forecast returns at longer horizons. To this end, we first use an exponential moving average to sum order flows into the past. We then use these lower frequency flows to build BMS portfolios that we rebalance every 2, 3, 4, 5, 10, 20, and 60 trading days. We report results for two different decay parameters (0.25 and 0.75) in the exponential moving average in Table IA.XII. We find that predictability dies out fairly quickly, although long-term demand-side investment managers' flows have some predictive power over longer horizons of up to one month (20 trading days).

Liquidity effects. To rule out the possibility that a simple liquidity story drives our predictability results, we also look at the subsample of the four most liquid currency pairs in our sample: EUR/USD, JPY/USD, GBP/USD, and CHF/USD. Table IA.XIII reports results for BMS portfolio returns and Figure IA.2 shows results for BMS returns in event time (similar to Figure 2 in the main text). We find that our main results remain qualita-

tively unchanged.

Individual currencies. We next explore whether a specific currency is driving the profitability of the order flow portfolios. To investigate this question, we rely on a cross-validation setting in which we form portfolios as before but in each case delete one of the available currencies. For example, we exclude the EUR/USD pair and compute BMS portfolio returns for the remaining 14 (total order flows) or eight currency pairs (disaggregated order flows). Table IA.XIV summarizes the results from this exercise. We continue to find the same general return pattern, which suggest that our main findings do not depend strongly on any particular currency.

In addition, we investigate trading strategies based on individual currencies where we go long (short) on a currency whenever the order flow on the previous day is positive (negative). As can be seen from the results reported in Table IA.XV, long-term demand-side investment managers do relatively well on almost all currencies (except CAD, NOK), whereas short-term demand-side investment managers tend to perform well only for the Scandinavian currencies SEK, NOK, and, to a lesser extent, NZD. Commercial corporations' performance is quite mixed, as expected, and individual investors show generally negative performance.²⁴

Order flow and macro fundamentals. Finally, we examine whether order flows are related to future macro fundamentals as suggested, for example, by Evans and Lyons (2008). We investigate this question in a cross-sectional setting, focusing on long-term demand-side investment managers' flows because we find in Figure 1 and Table IA.XII that long-term demand-side investment managers are the only end-users with permanent

forecasting ability. Figure IA.4 shows results for a simple exercise in which we forecast real industrial production (IP) growth and CPI inflation differentials based on lagged long-term demand-side investment managers' order flows. More specifically, the figure shows cumulative real IP growth and CPI inflation differentials for the group of countries in the BMS portfolio over time, that is, countries in Portfolio 4 minus countries in Portfolio 1. We employ a frequency of one month to match the availability of CPI and IP data. Hence, the figure illustrates the growth differentials between countries for which long-term demand-side investment managers exhibit the most intensive buying or selling pressure one period before. One can think of this as a cross-sectional out-of-sample test of predictability analogous to the portfolio sorts in Table I. Here, however, we look at growth rates in macroeconomic fundamentals and not currency returns.

Long-term demand-side investment managers' flows do indeed have sensible forecasting power for macroeconomic fundamentals, as shown in Figure IA.4. Their FX flows forecast higher growth in industrial output (p -value: 0.09) and lower inflation (p -value: 0.01), much in line with economic intuition. In other words, long-term demand-side investment managers overweight the currencies of countries with improvements in macroeconomic fundamentals relative to the currencies of the countries they underweight. While this exercise is intentionally simple, these findings are consistent with the notion of fundamental information processing by long-term demand-side investment managers, which helps explain why their trades have a permanent impact on exchange rates.

V. Conclusion

In this paper we empirically examine three related questions in an effort to improve our understanding of the ecology of the world’s largest financial market, the FX market. First, given that the FX market is fairly opaque and highly concentrated, how informative is observing a large proportion of the market’s order flow? Second, do FX end-users share risks among themselves, or is their trading highly correlated and unloaded onto the dealers and the interdealer market? Third, how can we understand the trading behavior, trading styles, and risk exposures of various key players in FX markets, and how is this linked to risk sharing?

We find that observing customer order flows is highly informative. Currency excess returns to portfolios mimicking aggregate customer order flows in real time are about 10% p.a. and highly significant. In addition, customer types vary massively in terms of their predictive ability, which matters especially because the FX market (like other OTC markets) is characterized by non-anonymity. Incorporating this feature into our setup, we find excess returns as high as 15% p.a., that is, non-anonymity further increases the informational value of order flow. The flows by long-term demand-side investment managers have the strongest predictive power for exchange rates, likely reflecting the ability to process fundamental information. Their flows have permanent forecasting power, whereas flows originating from the other groups only predict transitory changes in exchange rates.

We also find that the main segments of end-users differ markedly in their trading strategies and hedging demands. Moreover, flows of different end-user segments tend to be

negatively correlated over longer horizons. These findings suggest that risk sharing among end-users takes place not only via the interdealer market, as suggested by previous FX microstructure research, but also via the intermediation of large dealers.

These findings about information asymmetries, incentives, and risk sharing should be useful to inform policy discussions on the appropriate framework for OTC markets. Taken together, these results shed some light on one of the main OTC financial markets. Our findings suggest that the FX market is populated by quite heterogeneous market participants, and that we can gain valuable insights from observing their transactions and learning about their different predictive ability, trading motives, trading styles, and risk exposures.

Endnote

Initial submission: May 18, 2012; Final version received: July 31, 2015

Previous titles of paper: “*The Cross-Section of Currency Order Flow Portfolios*” and

“*Information Flows in Dark Markets: Dissecting Customer Currency Trades*”

Editor: Campbell R. Harvey

Notes

¹Lustig and Verdelhan (2007) were the first to build cross-sections of currency portfolios.

²See, for example, Payne (2003), Love and Payne (2003), Evans and Lyons (2002, 2008), Evans (2010), and Rime, Sarno, and Sojli (2010). Other papers relate order flow in a structural way to volatility (Berger, Chaboud, and Hjalmarsson (2009)) or directly to exchange rate fundamentals (Chinn and Moore (2011)).

³Several studies explore the underlying mechanism for the impact of order flow and discuss the evidence in terms of information versus liquidity effects (e.g., Berger et al. (2008), Cerrato, Sarantis, and Saunders (2011), Osler, Mende, and Menkhoff (2011), Menkhoff and Schmeling (2010), Phylaktis and Chen (2010), Moore and Payne (2011), Ito, Lyons, and Melvin (1998)).

⁴This information processing can manifest in various ways, for example, as more accurate and/or faster interpretation of macroeconomic news releases or better forecasting of market fundamentals such as the liquidity and hedging demands of other market participants.

⁵There is also evidence that marketwide private information extracted from equity order flow is useful for forecasting currency returns (Albuquerque, de Francisco, and Marques (2008)).

⁶Lustig and Verdelhan (2007), Farhi et al. (2013), Ang and Chen (2010), Burnside et al. (2011), Lustig, Roussanov, and Verdelhan (2011), Barroso and Santa-Clara (2015), and Menkhoff et al. (2012a, 2012b) all build currency portfolios to study return predictability and/or currency risk exposure.

⁷In addition, the volatility of flows varies over time and flows tend to become increasingly volatile towards the end of the sample. These features further call for some form of standardization.

⁸The nine currencies are: AUD, CAD, EUR, JPY, SEK, NZD, NOK, CHF, and GBP.

⁹It is important to note that we do not have data on individual customers and hence cannot use any

information on customers' identities; we only have data on customer types.

¹⁰This approximation is exact if covered interest rate parity (CIP) holds, which tends to be the case at daily or even shorter horizons in normal times (Akram, Rime, and Sarno (2008)). There have been violations of this no-arbitrage relation over the recent financial crisis. As we show below, the results in this paper are driven entirely by changes in spot rates, whereas interest rate differentials play only a negligible role. Thus, the results do not depend on whether CIP holds.

¹¹In these robustness exercises, we also report results with longer rolling windows of up to three years as well as for an expanding window. Furthermore, we conduct tests where we standardize both with respect to volatility as well as the mean. Finally, we consider a standardization scheme based on gross FX turnover data for different currencies drawing on data from the BIS FX triennial survey. These tests, reported in a separate Internet Appendix to conserve space, show that our results are not sensitive to the way in which flows are standardized. The Internet Appendix is available in the online version of the article on the Journal of Finance website.

¹²Subsample tests for a pre-crisis subperiod from January 2001 to June 2007 and a crisis/post-crisis subperiod from July 2007 to May 2011 are reported in the Internet Appendix.

¹³The MR statistic tests for a *monotonically* increasing return pattern, whereas the Up (Down) test is somewhat less restrictive and simply tests for a generally increasing (decreasing) pattern without requiring monotonicity in average portfolio returns. Specifically, the MR test requires that the return pattern be monotonically increasing $P_1 < P_2 < \dots < P_5$ and formulates the null hypothesis as $H_0 : \Delta \leq 0$ and the alternative hypothesis as $H_a : \min_{i=1,\dots,4} \Delta_i > 0$, where Δ is a vector of differences in adjacent average portfolio excess returns $(P_2 - P_1, P_3 - P_2, P_4 - P_3, P_5 - P_4)$ and Δ_i is element i of this vector. The Up test formulates the null hypothesis of a flat pattern $H_0 : \Delta = 0$ and the alternative hypothesis as $H_a : \sum_{n=1}^4 |\Delta_i| \mathbf{1}\{\Delta_i > 0\} > 0$, and hence it is less restrictive and also takes into account the size and magnitude of deviations from a flat return pattern. The Down test is constructed analogously.

¹⁴Table IA.IV in the Internet Appendix reports results for spot rate changes instead of excess returns,

which display no qualitative differences.

¹⁵Table IA.V in the Internet Appendix also shows that excess returns to the BMS portfolios based on different customers' flows are not highly correlated. Hence, the information contained in the different flows appears to stem from different sources. In practice, this also means that BMS portfolios could be combined to obtain even higher Sharpe Ratios. For example, a combined portfolio long in the long-term demand-side investment managers' BMS portfolio and short in the individual investors' BMS portfolio yields an annualized Sharpe Ratio of 2.19, which is substantially higher than the individual Sharpe Ratios.

¹⁶One strand of literature argues that order flow is the conduit by which information about fundamentals is impounded into prices and therefore has a permanent effect on exchange rates (e.g., Evans and Lyons (2002), Brandt and Kavajecz (2004), Evans and Lyons (2008)). Another strand of the literature suggests that order flow matters due to downward-sloping demand curves or "illiquidity," and hence order flow has only a transitory impact on prices (Froot and Ramadorai (2005)).

¹⁷This may be interpreted in the context of the adaptive markets hypothesis (see, for example, Neely, Weller, and Ulrich (2009) for an analysis in FX markets).

¹⁸Using more than one lag of flows in the regressions generally yields insignificant coefficient estimates so we restrict the regressions to include one lag of flows.

¹⁹Additional evidence is provided in the Internet Appendix. Table IA.XVIII summarizes exposures to equity factors, Table IA.XIX considers FX factors, Table IA.XX focuses on the Fung and Hsieh (2002) factors, and Table IA.XXI reports results for the BMS portfolio based on total flows for completeness.

²⁰Table IA.XVII reports pricing errors for the cross-section of order flow portfolios. Specifically, we report the Gibbons, Ross, and Shanken (1989) test for the null that the alphas are jointly equal to zero. Corroborating the time-series regressions in Tables IV and V, the test always rejects the null of zero alphas.

²¹The bid-ask spread data are available for *quoted spreads* and not effective spreads. As it is known that quoted spreads are much higher than effective spreads, we follow earlier work, for example, Goyal

and Saretto (2009), and employ 50% of the quoted bid-ask spread as the actual spread. Even this number seems conservative though. First, banks with access to this kind of customer order flow data are big dealers and pay very low spreads since they are key market makers. Second, Gilmore and Hayashi (2011) find in a recent study that transaction costs due to bid-ask spreads are likely to be much lower than our 50% rule.

²²Using other windows of less or more than 60 trading days does not yield qualitatively different results.

²³We linearly interpolate data in the BIS survey to obtain a daily time-series of trading volumes in USD for the nine developed currencies and then use the ratio of customer flows to total trading volumes as our sorting variable.

²⁴We also report results for a portfolio of all individual trading strategies for each customer group (last column in the table), which is more comparable to our order flow portfolios above. Qualitatively, the results are very similar and we find positive returns for financial customers but negative returns for individual clients. However, because the individual trading strategies are not dollar-neutral, the correlation between these trading strategies' returns with the returns of our dollar-neutral, cross-sectional BMS portfolios in Table I above are quite low and often negative.

REFERENCES

- Akram, Q. Farooq, Dagfinn Rime, and Lucio Sarno, 2008, Arbitrage in the foreign exchange market: Turning on the microscope, *Journal of International Economics* 76, 237–253.
- Albuquerque, Rui, Eva de Francisco, and Luis Marques, 2008, Marketwide private information in stocks: Forecasting currency returns, *Journal of Finance* 63, 2297–2343.
- Ang, Andrew, and Joseph Chen, 2010, Yield curve predictors of foreign exchange returns, Working paper, Columbia University.
- Bank for International Settlements, 2013, Triennial central bank survey, foreign exchange turnover in April 2013: Preliminary global results.
- Barber, Brad M., Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean, 2009, Just how much do individual investors lose by trading? *Review of Financial Studies* 22, 609–632.
- Barroso, Pedro, and Pedro Santa-Clara, 2015, Beyond the carry trade: Optimal currency portfolios, *Journal of Financial and Quantitative Analysis* 50, 1037–1056.
- Berger, David B., Alain P. Chaboud, Sergey V. Chernenko, Edward Howorka, and Jonathan H. Wright, 2008, Order flow and exchange rate dynamics in electronic brokerage system data, *Journal of International Economics* 75, 93–109.
- Berger, David B., Alain P. Chaboud, and Erik Hjalmarsson, 2009, What drives volatility persistence in the foreign exchange market? *Journal of Financial Economics* 94, 192–213.
- Brandt, Michael W., and Kenneth A. Kavajecz, 2004, Price discovery in the U.S. Treasury market: The impact of orderflow and liquidity on the yield curve, *Journal of Finance* 59, 2623–2654.
- Brunnermeier, Markus, Stefan Nagel, and Lasse Pedersen, 2009, Carry trades and currency crashes, *NBER Macroeconomics Annual 2008* 23, 313–347.
- Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo, 2011, Do Peso problems explain the returns to the carry trade? *Review of Financial Studies* 24, 853–891.
- Cerrato, Mario, Nicholas Sarantis, and Alex Saunders, 2011, An investigation of customer order flow in the foreign exchange market, *Journal of Banking and Finance* 35, 1892–1906.
- Chinn, Menzie D., and Michael Moore, 2011, Order flow and the monetary model of exchange rates: Evidence from a novel data set, *Journal of Money, Credit and Banking* 43, 1599–1624.
- Evans, Martin D., 2010, Order flows and the exchange rate disconnect puzzle, *Journal of*

- International Economics* 80, 58–71.
- Evans, Martin D., and Richard K. Lyons, 2002, Order flow and exchange rate dynamics, *Journal of Political Economy* 110, 170–180.
- Evans, Martin D., and Richard K. Lyons, 2005, Meese-Rogoff redux: Micro-based exchange rate forecasting, *American Economic Review, Papers and Proceedings* 95, 405–414.
- Evans, Martin D., and Richard K. Lyons, 2007, Exchange rate fundamentals and order flow, Working paper, Georgetown University.
- Evans, Martin D., and Richard K. Lyons, 2008, How is macro news transmitted to exchange rates? *Journal of Financial Economics* 88, 26–50.
- Farhi, Emmanuel, Samuel P. Fraiburger, Xavier Gabaix, Romain Ranciere, and Adrien Verdelhan, 2013, Crash risk in currency markets, Working paper, MIT Sloan.
- Ferson, Wayne E., and Rudi W. Schadt, 1996, Measuring fund strategy and performance in changing economic conditions, *Journal of Finance* 51, 425–462.
- Froot, Kenneth A., and Tarun Ramadorai, 2005, Currency returns, intrinsic value, and institutional-investor flows, *Journal of Finance* 60, 1535–1566.
- Fung, William, and David A. Hsieh, 2001, The risk in hedge fund strategies: Theory and evidence from trend followers, *Review of Financial Studies* 14, 313–341.
- Fung, William, and David A. Hsieh, 2002, Asset-based style factors for hedge funds, *Financial Analysts Journal* 58, 16–27.
- Fung, William, and David A. Hsieh, 2004, Hedge fund benchmarks: A risk based approach, *Financial Analysts Journal* 60, 65–80.
- Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121–1152.
- Gilmore, Stephen, and Fumio Hayashi, 2011, Emerging market currency excess returns, *American Economic Journal: Macroeconomics* 3, 85–111.
- Goyal, Amit, and Alessio Saretto, 2009, Cross-section of option returns and volatility, *Journal of Financial Economics* 94, 310–326.
- Grinblatt, Mark, and Matti Keloharju, 2000, The investment behavior and performance of various investor types: A study of Finland’s unique data set, *Journal of Financial Economics* 55, 43–68.
- Hasbrouck, Joel, 1991a, Measuring the information content of stock trades, *Journal of Finance* 46, 179–207.
- Hasbrouck, Joel, 1991b, The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies* 4, 571–595.

- Hau, Harald, and Helene Rey, 2004, Can portfolio rebalancing explain the dynamics of equity returns, equity flows, and exchange rates? *American Economic Review, Papers and Proceedings* 94, 126–133.
- Hvidkjaer, Søren, 2008, Small trades and the cross-section of stock returns, *Review of Financial Studies* 21, 1–23.
- Ito, Takatoshi, Richard K. Lyons, and Michael T. Melvin, 1998, Is there private information in the FX market? *Journal of Finance* 53, 1111–1130.
- Kaniel, Ron, Gideon Saar, and Sheridan Titman, 2008, Individual investor trading and stock returns, *Journal of Finance* 63, 273–310.
- King, Michael R., Carol Osler, and Dagfinn Rime, 2012, Foreign exchange market structure, players and evolution, in Jessica James, Ian Marsh, and Lucio Sarno, eds.: *Handbook of Exchange Rates* (Wiley, Hoboken, NJ).
- Love, Ryan, and Richard Payne, 2008, Macroeconomic news, order flows, and exchange rates, *Journal of Financial and Quantitative Analysis* 43, 467–488.
- Lustig, Hanno, Nikolai L. Roussanov, and Adrien Verdelhan, 2011, Common risk factors in currency markets, *Review of Financial Studies* 24, 3731–3777.
- Lustig, Hanno, Nikolai L. Roussanov, and Adrien Verdelhan, 2014, Countercyclical currency risk premia, *Journal of Financial Economics* 111, 527–553.
- Lustig, Hanno, and Adrien Verdelhan, 2007, The cross section of foreign currency risk premia and consumption growth risk, *American Economic Review* 97, 89–117.
- Lyons, Richard K., 1997, A simultaneous trade model of the foreign exchange hot potato, *Journal of International Economics* 42, 275–298.
- Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012a, Carry trades and global foreign exchange volatility, *Journal of Finance* 67, 681–718.
- Menkhoff, Lukas, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf, 2012b, Currency momentum strategies, *Journal of Financial Economics* 106, 660–684.
- Menkhoff, Lukas, and Maik Schmeling, 2010, Whose trades convey information? Evidence from a cross-section of traders, *Journal of Financial Markets* 13, 101–128.
- Moore, Michael, and Richard Payne, 2011, On the existence of private information in FX markets, *Journal of Banking and Finance* 35, 1250–1262.
- Neely, Christopher, Paul Weller, and Joshua M. Ulrich, 2009, The adaptive markets hypothesis: Evidence from the foreign exchange market, *Journal of Financial and Quantitative Analysis* 44, 467–488.
- Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, het-

- eroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.
- Osler, Carol L., Alexander Mende, and Lukas Menkhoff, 2011, Price discovery in foreign exchange markets, *Journal of International Money and Finance* 30, 1696–1718.
- Patton, Andrew J., and Tarun Ramadorai, 2013, On the high-frequency dynamics of hedge fund risk exposures, *Journal of Finance* 68, 597–635.
- Patton, Andrew J., and Allan Timmermann, 2010, Monotonicity in asset returns: New tests with applications to the term structure, the CAPM, and portfolio sorts, *Journal of Financial Economics* 98, 605–625.
- Payne, Richard, 2003, Informed trade in spot foreign exchange markets: An empirical investigation, *Journal of International Economics* 61, 307–329.
- Phylaktis, Kate, and Long Chen, 2010, Asymmetric information, price discovery and macroeconomic announcements in FX market: Do top trading banks know more? *International Journal of Finance and Economics* 15, 228–246.
- Rime, Dagfinn, Lucio Sarno, and Elvira Sojli, 2010, Exchange rates, order flow and macroeconomic information, *Journal of International Economics* 80, 72–88.
- Sager, Michael, and Mark P. Taylor, 2008, Commercially available order flow data and exchange rate movements: Caveat emptor, *Journal of Money, Credit and Banking* 40, 583–625.
- Whaley, Robert E., 2000, The investor fear gauge, *Journal of Portfolio Management* 26, 12–17.

Table I
Order Flow Portfolios: Excess Returns

This table reports average annualized portfolio excess returns for currency portfolios sorted on lagged order flow. We standardize order flow over a rolling window of 60 trading days prior to the order flow signal as outlined in the text. Column “Av” shows average excess returns across all currencies. Column “BMS” (bought minus sold) reports average excess returns for long-short portfolios in currencies with the highest versus lowest order flow. Numbers in brackets are t -statistics based on Newey-West standard errors whereas numbers in parentheses show (annualized) Sharpe Ratios. Columns “MR”, “Up”, and “Down” report p -values for tests of return monotonicity. The frequency is daily and the sample is from January 2001 to May 2011. Panel A reports results for total order flows and all 15 markets (T15) as well as for total order flows and the subsample of nine developed markets (T9). Panel B reports results for order flows disaggregated by customer type: long-term demand-side investment managers (LT), short-term demand-side investment managers (ST), commercial corporations (CO), and individual investors (II).

Panel A. Total Order Flows										
	P_1	P_2	P_3	P_4	P_5	Av.	BMS	MR	Up	Down
T15	0.82 [0.29] (0.09)	1.05 [0.37] (0.11)	6.15 [2.23] (0.71)	6.77 [2.40] (0.77)	11.13 [4.04] (1.21)	5.18 [2.20] (0.69)	10.31 [4.05] (1.26)	0.00	0.00	–
T9	0.34 [0.10] (0.03)	2.24 [0.74] (0.23)	8.21 [2.60] (0.80)	12.76 [4.17] (1.23)		5.89 [2.15] (0.66)	12.43 [4.68] (1.45)	0.00	0.00	–
Panel B. Disaggregated Order Flows										
LT*	-1.13 [-0.35] (-0.11)	3.75 [1.24] (0.38)	6.30 [2.04] (0.62)	14.31 [4.63] (1.38)			15.43 [5.72] (1.79)	0.00	0.00	–
ST*	-0.32 [-0.10] (-0.03)	6.05 [2.04] (0.61)	6.26 [1.94] (0.59)	9.78 [3.02] (0.94)			10.09 [3.94] (1.20)	0.04	0.00	–
CO	6.90 [2.15] (0.67)	5.27 [1.73] (0.53)	7.02 [2.16] (0.66)	2.61 [0.84] (0.26)			-4.29 [-1.66] (-0.51)	0.35	–	0.09
II	12.71 [4.06] (1.23)	6.69 [2.18] (0.67)	2.90 [0.93] (0.28)	-1.30 [-0.41] (-0.13)			-14.01 [-5.20] (-1.55)	0.00	–	0.00

*Long-term demand-side investment managers comprise “real money investors,” such as mutual funds and pension funds, whereas short-term demand-side investment managers comprise other funds and proprietary trading firms.

Table II
Order Flow Portfolios: Marginal Forecast Performance for Longer Horizons

This table reports average excess returns (p.a.) for BMS portfolios sorted on lagged order flow as in Table I. *t*-statistics based on Newey-West standard errors are reported in brackets. We not only sort on order flow of the previous day but also allow for longer lags of up to nine days between order flow signals and portfolio formation. Portfolios are rebalanced daily. T15 denotes portfolio sorts on total order flows and the sample of all 15 currencies. T9 denotes portfolio sorts on total order flows and the sample of nine developed currencies. LT, ST, CO, and II denote portfolio sorts on long-term demand-side investment managers', short-term demand-side investment managers', commercial corporations', and individual investors' order flows, respectively.

	Lags between order flow signal and portfolio formation (days)									
	1	2	3	4	5	6	7	8	9	10
T15	10.31 [4.05]	24.63 [8.94]	10.22 [4.38]	-1.11 [-0.44]	3.02 [1.28]	0.20 [0.09]	0.31 [0.13]	1.93 [0.84]	-2.32 [-0.95]	-0.43 [-0.19]
T9	12.43 [4.68]	24.27 [8.73]	7.44 [2.99]	-4.17 [-1.61]	5.39 [2.00]	-1.55 [-0.61]	2.28 [0.90]	1.33 [0.51]	-1.08 [-0.42]	-1.75 [-0.71]
LT*	15.43 [5.72]	24.86 [8.80]	8.27 [3.03]	-1.29 [-0.47]	2.17 [0.87]	0.62 [0.23]	-0.20 [-0.07]	3.37 [1.22]	2.26 [0.82]	-2.79 [-0.97]
ST*	10.09 [3.94]	28.22 [9.26]	2.05 [0.79]	-2.94 [-1.15]	0.14 [0.05]	-6.19 [-2.39]	2.84 [1.12]	-0.29 [-0.10]	-4.66 [-1.77]	-1.05 [-0.40]
CO	-4.29 [-1.66]	-8.13 [-2.86]	-1.47 [-0.49]	2.25 [0.88]	-4.98 [-1.93]	1.91 [0.74]	-0.01 [0.00]	1.40 [0.56]	-0.33 [-0.12]	2.80 [1.08]
II	-14.01 [-5.20]	-33.77 [-10.80]	3.21 [1.24]	1.82 [0.67]	-3.29 [-1.15]	-0.77 [-0.27]	2.27 [0.86]	-1.35 [-0.52]	0.65 [0.24]	2.10 [0.78]

*Long-term demand-side investment managers comprise "real money investors," such as mutual funds and pension funds, whereas short-term demand-side investment managers comprise other funds and proprietary trading firms.

Table III
Drivers of Customer FX Order Flow: Panel Regressions

This table reports results for panel regressions of customer order flows (OF) on lagged customer order flow (OF_t for long-term demand-side investment managers, LT, short-term demand-side investment managers, ST, commercial corporations, CO, and individual investors, II). The regressions also consider lagged returns on various asset classes as additional regressors (the interest rate differential $i_{j,t}^* - i_t$, lagged exchange rate changes over the previous day Δs_t and over the prior 20 trading days $\Delta s_{t-1,t-20}$, and lagged country-level equity returns over the previous trading day r_t^{eq} and over the prior 20 trading days $r_{t-1;t-20}^{eq}$), and lagged country-level government bond returns r_t^b (10-year maturity benchmark bonds). t -statistics based on clustered standard errors (by currency pair) are reported in brackets and we account for currency pair and time fixed effects.

	Dependent variable: Customer order flows							
	OF_{t+1}^{LT*}	OF_{t+1}^{ST*}	OF_{t+1}^{CO}	OF_{t+1}^{II}				
OF_t^{LT}	0.035 [4.46]	0.033 [4.22]	0.013 [1.79]	0.012 [1.73]	-0.010 [-1.13]	-0.009 [-1.08]	-0.005 [-0.61]	-0.003 [-0.38]
OF_t^{ST}	0.034 [2.75]	0.031 [2.66]	0.008 [0.57]	0.007 [0.50]	-0.009 [-1.70]	-0.008 [-1.55]	-0.037 [-2.59]	-0.350 [-2.56]
OF_t^{CO}	-0.017 [-2.58]	-0.016 [-2.53]	0.000 [0.02]	0.000 [0.05]	0.035 [2.93]	0.034 [2.88]	-0.012 [-1.47]	-0.013 [-1.62]
OF_t^{II}	-0.026 [-2.10]	-0.025 [-2.05]	-0.005 [-0.67]	-0.004 [-0.61]	0.025 [2.47]	0.024 [2.46]	0.027 [2.21]	0.025 [2.02]
$i_{j,t}^* - i_t$		-0.150 [-2.05]		0.102 [0.80]		0.413 [2.51]		0.185 [1.02]
Δs_t		3.541 [4.97]		1.769 [1.47]		-1.312 [-1.15]		-4.187 [-2.52]
$\Delta s_{t-1,t-20}$		1.012 [1.97]		0.612 [0.50]		-0.741 [-0.45]		-2.187 [-1.52]
r_t^{eq}		1.251 [2.56]		0.399 [0.41]		-1.164 [-2.37]		-0.226 [-0.34]
$r_{t-1;t-20}^{eq}$		0.347 [1.44]		-0.113 [-0.52]		0.205 [1.17]		-0.225 [-0.34]
r_t^b		-3.730 [-1.56]		-5.170 [-1.26]		-1.135 [-0.57]		10.145 [2.68]
$r_{t-1;t-20}^b$		-0.019 [-0.03]		0.278 [-0.55]		0.626 [1.04]		1.151 [2.03]
const.	0.008 [0.71]	-0.002 [0.03]	-0.078 [-4.42]	-0.089 [-3.47]	-0.320 [-7.27]	-0.295 [-6.01]	0.039 [4.77]	0.076 [4.41]
Country dummies	YES	YES	YES	YES	YES	YES	YES	YES
Time dummies	YES	YES	YES	YES	YES	YES	YES	YES
R^2	0.013	0.015	0.011	0.011	0.029	0.030	0.015	0.018
obs	23,796	23,796	23,796	23,796	23,796	23,796	23,796	23,796

*Long-term demand-side investment managers comprise “real money investors,” such as mutual funds and pension funds, whereas short-term demand-side investment managers comprise other funds and proprietary trading firms.

Table IV
Risk Exposures of Investment Managers

This table reports regression results for the risk exposures of the BMS portfolios of financial FX market end-users, that is, long-term demand-side investment managers (LT) and short-term demand-side investment managers (ST). The methodological framework in Panel A is a modified linear Fung-Hsieh (2002, 2004) model with eight factors as outlined in the main text. Panel B also accounts for conditional equity market exposures by including additional interaction terms. The three conditioning variables are first differences of the TED spread, the VIX and the 3-month T-bill rate. t -statistics based on HAC standard errors are reported (in brackets).

Panel A. Linear exposures								
	LT*				ST*			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
$PTFS_{FX}$		2.35 [2.65]				-2.68 [-2.51]		
$PTFS_{IR}$			3.07 [4.03]	2.18 [2.86]			-1.33 [-1.85]	-1.16 [-1.67]
ΔTS			-2.03 [-2.06]				0.38 [0.59]	
ΔDF			3.15 [2.83]	3.65 [2.87]			-3.58 [-2.61]	-3.67 [-2.69]
VOL_{FX}	0.07 [2.44]	0.06 [2.13]			-0.07 [-2.50]	-0.05 [-2.09]		
$\hat{\alpha}$	1.46 [5.32]	1.40 [5.45]	1.26 [5.68]	1.23 [5.25]	0.71 [3.10]	0.78 [3.49]	0.89 [4.01]	0.90 [3.97]
\bar{R}^2	0.10	0.12	0.21	0.15	0.11	0.14	0.10	0.10
Sys-BIC	3.53	3.53	3.54	3.54	3.53	3.53	3.54	3.54
Panel B. Interaction terms								
	LT*				ST*			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
$r_m \cdot \Delta VIX(t-1)$	-0.25 [-2.95]	-0.26 [-2.69]	-0.26 [-2.47]		0.15 [2.30]	0.17 [3.08]	0.16 [3.32]	
$r_m \cdot \Delta TED(t-1)$	-0.18 [-2.44]	-0.19 [-2.82]	-0.20 [-3.06]	-0.20 [-2.48]	0.32 [4.14]	0.37 [5.74]	0.33 [4.58]	0.34 [4.63]
$PTFS_{FX}$			2.52 [2.31]				-2.54 [-2.48]	
$PTFS_{IR}$		2.11 [2.79]				-0.87 [-1.35]		
VOL_{FX}	0.05 [2.08]			0.06 [2.15]	-0.05 [-2.10]			-0.05 [-2.29]
$\hat{\alpha}$	1.35 [5.57]	1.18 [5.55]	1.20 [5.38]	1.44 [5.31]	0.72 [3.50]	0.85 [3.89]	0.86 [4.05]	0.67 [3.21]
\bar{R}^2	0.17	0.18	0.15	0.12	0.21	0.18	0.21	0.19
Sys-BIC	3.45	3.46	3.47	3.47	3.45	3.46	3.47	3.47

*Long-term demand-side investment managers comprise “real money investors,” such as mutual funds and pension funds, whereas short-term demand-side investment managers comprise other funds and proprietary trading firms.

Table V
Risk Exposures of Commercial Corporations and Individual Investors

This table reports regression results for the risk exposures of the BMS portfolios computed from the flows of commercial corporations (CO) or individual investors (II). The methodological framework in Panel A is a modified linear Fung-Hsieh (2002, 2004) model with eight factors as outlined in the main text. Panel B also accounts for conditional equity market exposures by including additional interaction terms. The three conditioning variables are first differences of the TED spread, the VIX, and the three-month T-Bill rate. Below the regression coefficients, t -statistics based on Newey and West standard errors are reported in brackets.

Panel A. Linear exposures								
	CO				II			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
r_m	-0.14 [-2.32]	-0.11 [-1.63]		-0.14 [-1.94]	-0.07 [-1.56]	-0.09 [-2.27]		-0.10 [-2.43]
$PTFS_{IR}$				-1.83 [-1.04]				-0.70 [-0.52]
ΔDF	-3.43 [-1.38]		-2.20 [-0.99]		2.57 [2.81]		3.18 [3.51]	
$\hat{\alpha}$	-0.30 [-1.49]	-0.31 [-1.5]	-0.37 [-1.73]	-0.25 [-1.42]	-1.16 [-4.27]	-1.15 [-4.13]	-1.19 [-4.34]	-1.12 [-4.04]
\bar{R}^2	0.07	0.04	0.01	0.07	0.06	0.03	0.04	0.03
Sys-BIC	3.93	3.94	3.94	3.96	3.93	3.94	3.94	3.96
Panel B. Interaction terms								
	CO				II			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
r_m		-0.08 [-1.30]	-0.10 [-1.50]			-0.11 [-2.48]	-0.13 [-3.05]	
$r_m \cdot \Delta TED(t-1)$	0.47 [2.95]	0.44 [2.29]	0.40 [2.1]	0.49 [2.56]	-0.20 [-2.53]	-0.24 [-2.27]	-0.31 [-2.68]	-0.12 [-0.99]
$PTFS_{IR}$			-0.94 [-0.8]				-1.38 [-1.45]	
ΔDF				0.39 [0.25]				2.51 [1.71]
$\hat{\alpha}$	-0.46 [-2.04]	-0.41 [-1.91]	-0.37 [-1.85]	-0.46 [-2.02]	-1.18 [-4.25]	-1.12 [-4.27]	-1.06 [-4.20]	-1.19 [-4.46]
\bar{R}^2	0.13	0.14	0.15	0.12	0.02	0.07	0.09	0.04
Sys-BIC	3.82	3.83	3.87	3.88	3.82	3.83	3.87	3.88

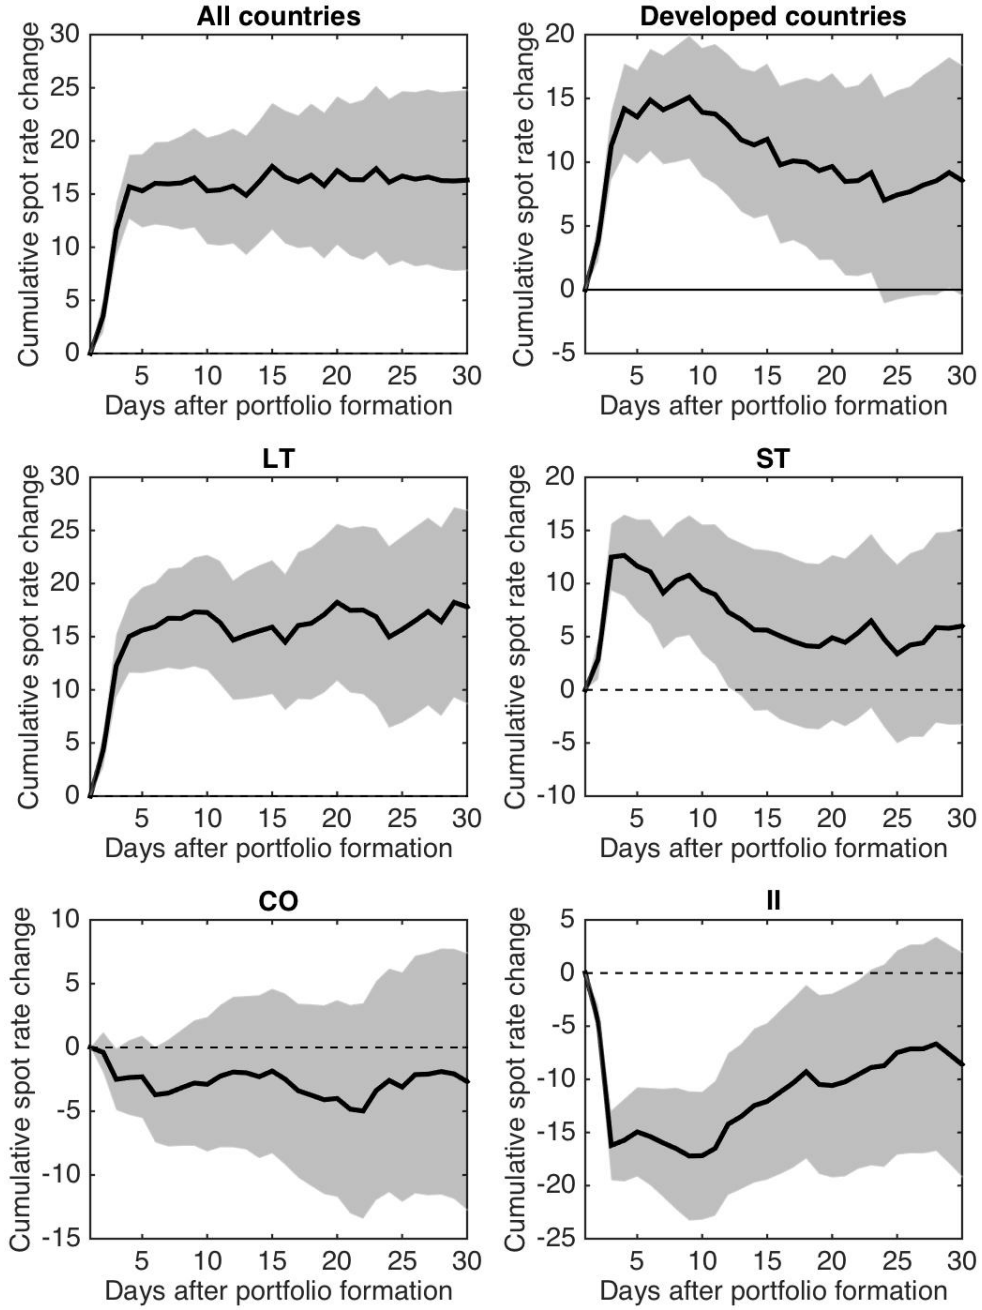


Figure 1. Cumulative post-formation exchange rate changes. This figure shows average cumulative spot exchange rate changes for BMS portfolios based on total flows and disaggregated flows over the first 30 days after portfolio formation. We use daily data so that post-formation periods overlap. LT denotes long-term demand-side investment managers, ST denotes short-term demand-side investment managers, CO denotes corporations, and II denotes individual investors. Long-term demand-side investment managers comprise “real money investors,” such as mutual funds and pension funds, whereas short-term demand-side investment managers comprise other funds and proprietary trading firms. Shaded areas correspond to a 95% confidence interval obtained from a moving-block bootstrap with 1,000 repetitions.




Figure 2. BMS excess returns in event time. This figure plots BMS portfolio excess returns (solid lines) in event time, from five days prior to portfolio formation ($t = -5$), the day of portfolio formation ($t = 0$), and up to 10 days after portfolio formation ($t = 10$). BMS excess returns are annualized and in %. Long-term demand-side investment managers comprise “real money investors,” such as mutual funds and pension funds, whereas short-term demand-side investment managers comprise other funds and proprietary trading firms. The frequency is daily and the sample is from January 2001 to May 2011.

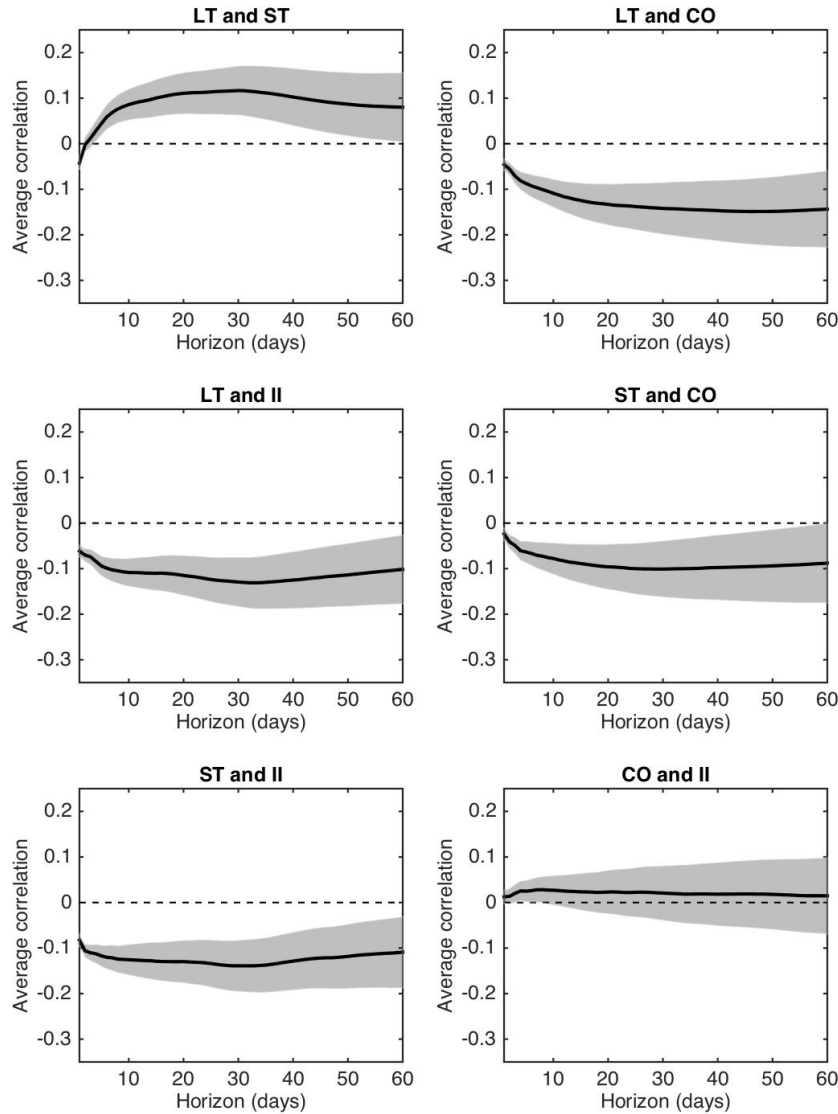


Figure 3. Correlation of customer order flows over longer horizons. This figure plots average correlation coefficients between customer order flows (left panel) for horizons of 1, 2, ..., and 60 trading days. Average correlations between flows are based on the average correlation across all nine currency pairs. A horizon of one day corresponds to (non-overlapping) daily observations, whereas correlations for longer horizons are based on (overlapping) sums of daily observations. Shaded areas correspond to bootstrapped 95% confidence intervals based on a moving-block bootstrap with 1,000 repetitions. LT denotes long-term demand-side investment managers, ST denotes short-term demand-side investment managers, CO denotes corporations, and II denotes individual investors. Long-term demand-side investment managers comprise “real money investors,” such as mutual funds and pension funds, whereas short-term demand-side investment managers comprise other funds and proprietary trading firms. The sample period is January 2001 to May 2011.