

Currency Order Flow and Real-Time Macroeconomic Information*

Pasquale DELLA CORTE

Imperial College London

p.dellacorte@imperial.ac.uk

Dagfinn RIME

Norges Bank

dagfinn.rime@norges-bank.no

Lucio SARNO

Cass Business School & CEPR

lucio.sarno@city.ac.uk

Ilias TSIAKAS

University of Guelph

itsiakas@uoguelph.ca

February 2013

***Acknowledgements:** The authors are indebted for constructive comments to Rui Albuquerque, Philippe Bacchetta, Ekkehart Boehmer, Nicola Borri, Giuseppe De Arcangelis, Martin Evans, Charles Jones, Nengjiu Ju, Michael King, Robert Kosowski, Michael Moore, Marco Pagano, Alessandro Palandri, Lasse Pedersen, Tarun Ramadorai, Thomas Stolper, Adrien Verdelhan, Paolo Vitale, Kathy Yuan, Sarah Zhang and seminar participants at LUISS Guido Carli University, University of Lugano, Warwick Business School, the 2011 Capital Markets and Corporate Finance Meetings in Kunming, the 2011 Central Bank Workshop on the Microstructure of Financial Markets in Stavanger, the 2011 Conference on Advances in the Analysis of Hedge Fund Strategies in London, the 2011 Workshop on Financial Determinants of Exchange Rates in Rome, the 2012 CFA Society Masterclass Series in London, the 2012 SIRE Econometrics Workshop in Glasgow, the 2012 Rimini Conference in Economics and Finance in Toronto, and the 2012 Northern Finance Association Conference in Niagara Falls. We thank UBS for providing the customer order flow data used in this paper. Sarno acknowledges financial support from the Economic and Social Research Council (No. RES-062-23-2340). Tsiakas acknowledges financial support from the Social Sciences and Humanities Research Council of Canada. The views expressed in this paper are those of the authors, and not necessarily those of Norges Bank. **Corresponding author:** Lucio Sarno, Cass Business School, City University, 106 Bunhill Row, London EC1Y 8TZ, UK. Email: lucio.sarno@city.ac.uk

Currency Order Flow and Real-Time Macroeconomic Information

Abstract

This paper investigates empirically whether currency order flow aggregates dispersed real-time macroeconomic information using a unique data set on customer order flow disaggregated across four customer groups for the G10 currencies over a ten-year sample period. We first establish that customer order flow has substantial out-of-sample forecasting ability for exchange rate returns in the context of a dynamic trading strategy with monthly rebalancing. We then find that a large part of the information in order flow can be explained ex post by a time-varying combination of real-time macroeconomic fundamentals. However, models conditioning on macroeconomic information fail to replicate ex ante models conditioning on order flow as the latter substantially outperform the former using economic metrics of forecast evaluation. This leads us to conclude that order flow provides a distinct and effective way of aggregating dispersed macroeconomic information.

Keywords: Exchange Rates; Order Flow; Real-Time Economic Fundamentals; Forecasting; Asset Allocation.

JEL Classification: F31; G11; G15.

1 Introduction

Does currency order flow aggregate dispersed macroeconomic information across different market participants? This question is at the center of recent theories of exchange rate determination developed by Bacchetta and van Wincoop (2004, 2006), building on the insights of the market microstructure approach to exchange rates (Evans and Lyons, 2002; Evans and Rime, 2012). This approach has emerged as an exciting alternative to traditional economic models of exchange rate determination, which despite thirty years of research have had limited success in explaining and predicting currency movements. As a result, exchange rates are thought to be largely disconnected from macroeconomic fundamentals. In contrast, it is a robust finding across currencies, sample periods and frequencies, that currency transactions affect exchange rates. The market microstructure literature asserts that this relation is due to information revealed by order flow.¹ In particular, macroeconomic news can be impounded directly in currency prices or indirectly via order flow (e.g., Albuquerque, de Francisco and Marques, 2008; Evans and Lyons, 2008; Osler, Mende and Menkhoff, 2011).² However, order flow can also affect prices for reasons unrelated to publicly available news, such as changing risk aversion, liquidity and hedging demands (e.g., Berger, Chaboud, Chernenko, Howorka and Wright, 2008; Breedon and Vitale, 2010).

This paper investigates empirically the relation between the predictive information in customer order flow and real-time macroeconomic information. The use of customer order flow data is especially important, as these data reflect the underlying motives for trade of heterogeneous customers who initiate trades with dealers acting as intermediaries. The use of real-time macroeconomic data is also important, as these data reflect the information actually available to customers at the time they initiate trades. In particular, we use a proprietary order flow data set obtained from UBS, a global leader in foreign exchange (FX) markets, on their daily trading with four customer segments: asset managers, hedge funds, corporates and private clients. This is a rich data set that contains the US dollar value of order flow for the period of January 2001 to May 2011 and covers the G10 currencies. The data set on macroeconomic variables is constructed using real-time data that was available to market participants when their trading decisions were made. The two data sets, therefore, provide us with a unique opportunity to examine the information content of customer order flow and its

¹Order flow is a measure of the net demand for a particular currency defined as the value of buyer-initiated orders minus the value of seller-initiated orders. Note that earlier studies use a simpler definition of order flow as the number (not value) of buyer-initiated trades minus the number of seller-initiated trades (e.g., Evans and Lyons, 2002).

²For example, consider a scheduled macro announcement that is better than expected by market participants. Suppose that everyone agrees that the announcement (e.g., on the current account) represents good news for the domestic currency but there are diverse opinions as to how large the appreciation should be. Those who view the initial rise as too small will place orders to purchase this currency, while those who view the initial rise as too large will place orders to sell. In aggregate, positive order flow signals that the initial spot rate was below the balance of opinion among market participants.

relation to real-time macroeconomic fundamentals over a long sample and a large set of exchange rates.

Armed with these data, our paper examines whether macroeconomic information can explain the predictive ability of order flow. We address this question in four steps. First, we establish whether there is valuable predictive information in currency order flow. To this end, we take the point of view of an investor (or dealer) implementing a dynamic asset allocation strategy across the G10 currencies. It is worth noting that the trading strategy endogenizes transaction costs so that the bid-ask spread directly affects the optimal weights. The trading strategy allows us to measure the tangible economic gains of conditioning out of sample on different types of customer order flow. It also sheds light on the trading decisions through which different customer groups reveal their information.

Second, we relate the realized portfolio returns generated by conditioning on customer order flow - a direct measure of the value of the information in order flow - to the realized portfolio returns generated by conditioning on the macroeconomic fundamentals commonly used in the literature. This way, we can assess both the extent to which customer trading decisions reflect, for example, changes in interest rates, real exchange rates or other economic fundamentals and the extent to which they reflect information not related to economic fundamentals. Applying a framework based on portfolio returns allows us to relate the value of the information in order flow to the value of macroeconomic information using the same units of measurement. The empirical analysis is carried out using monthly data as this is the frequency at which most macroeconomic information is released.

Third, we examine the extent to which the relation between order flow and macroeconomic information varies over time. This is an important question, since it is possible (even likely) that FX participants change over time the weight they assign to different fundamentals. This practice is consistent with the scapegoat theory of Bacchetta and van Wincoop (2004, 2011), where over time the market may focus its attention on a different macroeconomic variable (the scapegoat).³ The scapegoat theory relies on traders assigning a different weight to a macroeconomic indicator over time as the market rationally searches for an explanation for the observed exchange rate change.

Finally, we determine whether standard forecast combinations conditioning on macroeconomic fundamentals can replicate the out-of-sample forecasting ability of order flow in real time. If so, order flow does not make a meaningful contribution to exchange rate predictability in the sense that it simply combines widely available economic information in a manner that is straightforward to replicate. If not, it could be that order flow summarizes the available macroeconomic information in a distinct and effective manner that cannot be replicated in real time by simply combining forecasts based on public information.

³This practice is also documented in the survey evidence of Cheung and Chinn (2001) that is based on questionnaires sent to FX traders.

To put our empirical analysis in perspective, it is useful to summarize certain aspects of the trading mechanism for currencies. The FX market comprises two distinct groups of participants: dealers and end-user customers. Dealers act as financial intermediaries who facilitate trades by quoting prices at which they are willing to trade with customers. Trades between dealers and customers are not transparent, since there are no disclosure requirements in FX markets and, as a result, prices and transaction volumes are only observed by the two transacting counterparties. Therefore, customer orders are an important source of information to dealers as they may signal the customers' interpretation of public news, their needs and their required risk premia. This information is then revealed to the rest of the market when dealers trade with one another.⁴ Trades between dealers and customers account for 61.1% of FX turnover (Bank for International Settlements, 2010).⁵

This trading mechanism implies that customer order flow may contain information about future FX excess returns. The argument is as follows: Customer order flow is private information only available to the intermediating dealer. Customers are heterogeneous in their motivation for trading, beliefs, attitude towards risk and horizon. Therefore, different customers provide a dealer with different information. However, as dealers from different banks trade with one another the information in customer trades will eventually be reflected in the exchange rate. In short, the spot exchange rate is the rate quoted by FX dealers, and hence reflects the dealers' common information set, which in turn aggregates the information in customer order flow. Note that this learning process may take time due to the opaqueness of customer trades and the possibly delayed trading by dealers using customer information. As a result, in this mechanism, the information in customer order flow can generate not only a contemporaneous relation but also a predictive relation to FX excess returns (see, e.g., Evans and Lyons, 2005; Rime, Sarno and Sojli, 2010).

We find that customer order flow for currencies has substantial out-of-sample forecasting ability, which can lead to highly profitable trading strategies net of transaction costs. The order flow of asset managers and hedge funds tends to have the highest predictive ability, whereas the order flow of private clients is the least informative. Although the predictive ability of order flow is not the primary focus of our empirical analysis, it is worth noting that our out-of-sample forecasting results conditioning on order flow are far better than, for example, Evans and Lyons (2005) and Rime, Sarno and Sojli (2010). Furthermore, our order flow sample is longer and for more currencies.

⁴In the interdealer market, dealers have access to two different trading channels: they can trade directly with each other or through brokers, where the latter includes FX trading platforms such as Reuters and Electronic Broking System. The direct interdealer trades are private since the bid and ask quotes, the amount and direction of trade are not announced to the rest of the market. The second channel is more transparent as electronic brokers announce best bid and ask prices and the direction of all trades. However, this information is predominantly available to dealers (see, e.g., Bjonnes and Rime, 2005).

⁵For further details on the institutional structure of the FX market, see, for example, Lyons (2001), Sager and Taylor (2006), and King, Osler and Rime (2012).

Our main finding is that macroeconomic information can explain ex post a large part (up to 50%) of the predictive information of order flow. We summarize the relevant macroeconomic information using a number of empirical exchange rate models: the random walk, forward premium, uncovered interest parity, purchasing power parity, monetary fundamentals, Taylor rule, cyclical external imbalances, and momentum.⁶ Notably, there is no additional predictive information in order flow over and above the information embedded in macroeconomic fundamentals.

We also find that the relation between order flow and macroeconomic information can vary significantly over time as investors react to changes in their economic environment by assigning a different weight each month to macroeconomic fundamentals. This is particularly evident before versus after the crisis, as, for example, after the crisis the random walk (carry trade) has been replaced as an important driver of order flow by purchasing power parity, monetary fundamentals and the Taylor rule. This result is also consistent with the scapegoat theory of Bacchetta and van Wincoop (2004, 2011). By allowing for time-variation in the relation between order flow and macroeconomic information, the latter can explain 50% to 70% of the former, on average.

Finally, our analysis demonstrates that although order flow can be largely explained by macroeconomic information ex post, it is far from being redundant. Indeed, standard forecast combinations of macroeconomic fundamentals fail to replicate ex ante the out-of-sample forecasting ability of order flow. Order flow strategies can outperform the best forecast combination based on macroeconomic information by approximately 20% per year. In other words, if we knew ex ante the weights on the combination of macroeconomic variables, we would be able to replicate the predictive information in order flow. Since, however, we can only estimate these weights ex post in a regression that involves order flow data, the predictive ability of order flow cannot be replicated ex ante. This leads us to conclude that order flow provides a distinct and effective way of aggregating dispersed public information about economic fundamentals that are relevant to exchange rates, as predicted by microstructure theory.

Our analysis is related to a series of recent papers by Pojarliev and Levich (2008, 2010, 2012), who provide an empirical investigation of professionally managed currency funds. Their analysis identifies four risk factors that explain the returns to currency funds: (i) the carry factor based on the returns to the carry trade, (ii) the trend factor based on moving average rules, (iii) the value factor based

⁶Strictly speaking, the random walk and momentum strategies are not based on macroeconomic information but are useful in helping us understand what drives customer order flow. The random walk model has long been established as the benchmark in evaluating exchange rate predictability since Meese and Rogoff (1983). If the random walk holds, then FX spot returns cannot be predicted by macroeconomic information, but FX excess returns are perfectly predictable and equal to the interest rate differential. Since the returns to currency portfolio strategies depend on FX excess returns, then implicitly macroeconomic information has a prominent role under the random walk. On the other hand, momentum is not based on macroeconomic information but is a profitable and increasingly popular FX strategy (see Menkhoff, Sarno, Schmeling and Schrimpf, 2012b).

on purchasing power parity, and (iv) the volatility factor based on FX implied volatilities. A key difference between the analysis of Pojarliev and Levich (2008, 2010, 2012) and this paper is that they focus on explaining the observed returns of currency managers, whereas we study trading decisions based on the predictive information in order flow. Another related paper is Barroso and Santa-Clara (2012), who find that conditioning on the interest rate spread, momentum and reversal can generate optimal currency portfolios that outperform out-of-sample the carry trade net of transaction costs.

This paper is also related to recent work that attempts to bridge the macro-micro dichotomy of exchange rate determination by relating exchange rates, order flow and macroeconomic information. For example, Evans and Lyons (2008), Love and Payne (2008) and Rime, Sarno and Sojli (2010) use interdealer order flow data for up to three exchange rates, over a short sample of up to one year, at a frequency that is daily or higher. These papers study the relation between order flow and macroeconomic news announcements. Evans (2010) provides a similar empirical investigation with two main differences: he uses weekly customer order flows for one exchange rate over six years obtained from Citibank, and quantifies dealers' evolving views of macroeconomic conditions using an estimate of real-time expectation errors associated with each macroeconomic variable. These real-time expectation errors are not observed but are estimated using a subset of information that was available at the time based on monthly and quarterly releases of macroeconomic variables.

Our analysis deviates substantially from this previous work in a number of ways: (i) we use customer order flow disaggregated across four customer groups; (ii) our data sample comprises ten currencies for more than ten years; (iii) the frequency of the analysis is monthly, which is the frequency at which most macroeconomic information is released; (iv) we condition on a large set of observed real-time monthly macroeconomic fundamentals (rather than news releases) to capture the public information actually available at the time trading decisions are made; and (v) we explore the time-varying relation between currency order flow and macroeconomic information on the basis of the trading decisions through which customers reveal their information. Hence this paper provides a novel perspective on the extent to which currency order flow is an aggregation of dispersed macroeconomic information across different market participants.

The remainder of the paper is organized as follows. In the next section we describe the data used in the empirical analysis, with particular emphasis on the UBS data set for currency order flows and the data set on real-time macroeconomic variables. Section 3 describes the empirical models that link order flow, exchange rates and the macroeconomy. In Section 4 we present the dynamic asset allocation framework used to assess the predictability of customer order flow or macroeconomic information. Section 5 discusses the empirical results and, finally, Section 6 concludes.

2 Data and Preliminaries

The empirical analysis uses spot and forward exchange rates, customer order flows and a set of macroeconomic variables for nine exchange rates relative to the US dollar (USD): the Australian dollar (AUD), Canadian dollar (CAD), Swiss franc (CHF), euro (EUR), British pound (GBP), Japanese yen (JPY), Norwegian krone (NOK), New Zealand dollar (NZD) and Swedish krona (SEK). The data range from January 2001 to May 2011. Descriptive statistics for exchange rate returns and order flow are reported using daily data. However, our main empirical analysis that links order flow to macroeconomic variables is based on monthly data as this is the frequency at which most macroeconomic information is released.

2.1 Exchange Rate Data

We use end-of-month spot and one-month forward exchange rates obtained from Thomson Reuters through *Datastream*. The exchange rate is defined as the US dollar price of a unit of foreign currency so that an increase in the exchange rate implies a depreciation of the US dollar. The regressions that generate expected returns use mid-quotes, but bid and ask quotes are used to construct transaction costs and calculate the net returns of the trading strategies.

2.2 Order Flow Data

In a two-tiered market like the FX market, dealers are to a large extent merely intermediaries for customers, and it is primarily the customer transactions that carry information. Our customer order flow data come from proprietary daily transactions between UBS, a global leader in the FX market, and four end-user segments (customer groups): asset managers (AM), hedge funds (HF), corporates (CO) and private clients (PC). The asset managers segment comprises real money investors, such as mutual funds and pension funds. Highly leveraged traders and short-term asset managers not included in the asset managers segment are classified as hedge funds. Hedge funds are not regulated entities, whereas asset managers are. The corporate segment includes non-financial corporations that import or export products and services around the world or have an international supply chain. Corporates also include the treasury units of large non-financial corporations, with the exception of those pursuing a highly leveraged investment strategy, in which case they are classified as hedge funds. The last segment, private clients, includes wealthy clients with in excess of 3 million US dollars in investable liquid assets. Private clients trade primarily for financial reasons and with their own money.

Table 1 reports the market share of UBS by customer type and its rank relative to the top 10 global FX dealers from 2001 to 2011 based on the Euromoney annual survey. The table reveals that

UBS has been one of the top dealers for both the overall market and particular end-user segments. This is important as it is the large banks that determine the price discovery process in this highly concentrated market. The market relevance of the UBS customer groups can be seen from the close alignment between the AM, HF and CO segments of UBS and the real money, leveraged funds and non-financial corporations groups of the Euromoney survey. The table indicates that UBS is among the top two banks trading with asset managers, among the top five banks for hedge funds, and among the top ten banks for non-financial corporations. The latest estimate of the average daily turnover in FX is \$4.7 trillion (Bech, 2012). As UBS has an overall market share close to 10%, it approximately produces a daily turnover of more than \$400 billion per day, or more than \$8 trillion per month.⁷

The order flow data are assembled as follows. Each transaction booked in the UBS execution system at any of its world-wide offices is tagged with a client type. At the end of each business day, global transactions are aggregated for each customer group. Order flow is measured as the difference between the dollar value of purchase and sale orders for foreign currency initiated by a particular UBS customer group. Specifically, let V_t be the US dollar value of a transaction initiated by a customer at time t . The transaction is recorded with a positive sign if the initiator of the transaction (the non-quoting counterparty) is buying foreign currency and vice versa. It follows that positive order flow indicates a net demand of foreign currency, whereas negative order flow a net supply.⁸

The order flow data set used in our analysis is unique in many respects. First, in contrast to most other empirical studies that focus on interdealer data, we use customer order flow data disaggregated into the four segments discussed above. These customer segments are the result of an aggregation of all UBS transactions for each of the four types of customers. Decomposing UBS's total customer order flow into four aggregate measures is very useful in determining the predictive information content of order flow for different customers, and, more importantly, in assessing whether macroeconomic information can explain order flow information. Our prior is that the results will differ across different customer groups. In the end, we are not interested in the actual returns of a particular fund or customer but in the content and drivers of the predictive information of an aggregate customer group.

Second, our data set spans more than ten years with 2,616 daily observations for nine currency pairs and it comes from a major FX market leader. Although there are recent studies that employ

⁷Note that the Euromoney FX Survey also has a group called Banks, which covers so-called non-market making banks, often small banks, that do not find it worthwhile to have a presence in the interbank market but rather trade with other banks as their customer. There is no similar group in the UBS definitions. Furthermore, based on our private discussions with traders, UBS is also a leading counterparty trader with private clients.

⁸It is important to note that order flow is distinct from transaction volume. Order flow is transaction volume that is signed. Microstructure theory defines the sign of a trade depending on whether the initiator (i.e., customer) is buying or selling. Consider, for example, a sale of 10 units of foreign currency by a customer acting on a dealer's quotes. Then transaction volume is 10, but order flow is -10 (see, e.g., Lyons, 2001).

customer order flow data, they typically suffer from various limitations, as they cover a relatively short period of time, fewer currency pairs or a limited number of end-user segments. For instance, Evans and Lyons (2005) and Evans (2010) employ six years of data for one currency pair from Citibank. Froot and Ramadorai (2005) use seven years of data for eighteen currency pairs from State Street, a global custodian bank. These are flow data with primarily institutional investors, which are however aggregated, and hence cannot capture the same diversity in currency demand as the UBS end-user segments.

Third, many empirical studies use the number (not the dollar value) of buyer-initiated and seller-initiated transactions to measure order flow (e.g., Evans and Lyons, 2002). Particularly for customer order flow, it is important to use a measure based on the dollar value of transactions because there can be large differences in the value of trades across different customer groups. For example, the trades of private clients may on average be much smaller than the trades of asset managers, and our measure of order flow accounts for that.

Fourth, our order flow data are raw data with filtering limited to the most obvious cases. For instance, data are adjusted to take into account large merger and acquisition deals which are announced weeks or months in advance. Cross-border mergers and acquisitions involve large purchases of foreign currency by the acquiring company to pay the cash component of the deal. These transactions are generally well-publicized and thus are anticipated in advance by market participants. Finally, FX reserve managers, UBS proprietary (prop) traders and small banks not participating in the interbank market are excluded from the data set. Flows from FX reserve managers are stripped out due to confidentiality issues, flows from prop traders because they trade with UBS' own money, while small banks represent small customers less concerned about the FX market.

2.3 Real-Time Macroeconomic Data

Macroeconomic data typically used in the literature are not real-time data as they are subject to a number of revisions when more accurate estimates become available. Hence, such final-release macroeconomic data cannot be used to replicate the information available to investors in real time. In contrast, real-time data refer to vintage versions of economic data that were available on a given date in history. An important advantage of our empirical analysis is that we manually collect and compile real-time economic data from historical paper and electronic sources for the following countries: Australia, Canada, the euro area, Japan, New Zealand, Norway, Sweden, Switzerland, the United Kingdom, and the United States. We construct monthly data on industrial production, real output, consumer prices, money supply, external assets and liabilities, and exports and imports. Starting from December 2003, we construct twelve vintages per year for a total of 90 real-time vintages until May 2011. Each vintage goes back in time to January 2001. When data are only available at

lower frequency, we retrieve monthly data by linear interpolation and linear extrapolation.⁹ Seasonal adjustment is performed by the relevant statistical offices in real time. Overall, the real-time data is assembled in a way that avoids forward-looking bias. This allows us to conduct an accurate real-time analysis, where time t information refers to the information actually available at time t for each macroeconomic variable. A list of the data sources is reported below.

We obtain real-time seasonally adjusted data on industrial production, GDP, consumer price indices, and broad money from the OECD *Main Economic Indicators*. Industrial production indices are available at monthly frequency, except for Australia, New Zealand and Switzerland which are quarterly. The GDP data are available in national currency at quarterly frequency for all countries. Consumer prices are published every month for all countries, except Australia and New Zealand for which they are available quarterly. Broad money refers to the monetary aggregate M3 for all countries, except Canada (M2), Japan (M4), Norway (M2), and the United Kingdom (M4). These data are available in national currency at monthly frequency.

For foreign assets and liabilities, we collect real-time data on external assets and liabilities from the IMF *International Financial Statistics*. These data are published in US dollars every quarter for most of the countries. For Japan, Norway, and the United States data are only available at annual frequency. For exports and imports, we collect real-time seasonally adjusted data on exports and imports of goods and services from the OECD *Quarterly National Accounts*. These data are published every quarter in national currency for all countries. Finally, for monthly interest rates, we use end-of-month Eurodeposit rates from *Datastream*.

2.4 Preliminary Analysis

Table 2 presents descriptive statistics for daily log exchange rate returns and order flows across the four customer groups for the nine US dollar exchange rates from January 2001 to May 2011. Order flow tends to be more volatile for asset managers and hedge funds and least volatile for corporates. This fits with the view that asset managers and hedge funds are active traders, whereas corporate clients trade mostly for import and export reasons.

Table 3 presents the contemporaneous cross-correlations between daily customer order flows and exchange rate returns. While the order flows of asset managers and hedge funds are positively correlated with exchange rate returns, the order flows of corporates and private clients tend to be negatively correlated. Furthermore, the order flows of asset managers and hedge funds tend to be significantly negatively correlated with the order flows of corporates and private clients. This is an

⁹For example, on March 2004 data on Australian exports were only available at quarterly frequency until September 2003 (hence there was no data for the period of October 2003 to March 2004). We use linear interpolation to construct monthly observations up to September 2003, and linear extrapolation for the missing monthly observations up to March 2004.

interesting preliminary finding, which indicates that different types of order flow may move exchange rates in opposite directions. These results are also consistent with the framework of Evans and Lyons (2002) and Sager and Taylor (2006), who argue that asset managers and hedge funds are informed traders (push customers) whereas corporate and private clients act as overnight liquidity providers (pull customers).

3 Predictive Regressions

This section describes two sets of predictive regressions for exchange rate returns. The first set conditions on customer order flows, whereas the second set conditions on different macroeconomic fundamentals suggested by exchange rate theories. All predictive regressions have the following linear structure:

$$\Delta s_{t+1} = a + bx_t + \varepsilon_{t+1}, \quad (1)$$

where s_{t+1} is the nominal US dollar spot exchange rate for a particular currency at time $t + 1$, $\Delta s_{t+1} = s_{t+1} - s_t$ is the log-exchange rate return at time $t + 1$, x_t is a predictive variable, a and b are constant parameters to be estimated, and ε_{t+1} is a normal error term. The regressions differ only in the way they specify the predictive variable x_t that is used to forecast exchange rate returns.

For the macroeconomic models we constrain the sign of the slope parameter b . This is important, and in fact necessary, for the predictive regression to represent the economic theory behind it. For example, in the predictive regression based on purchasing power parity, we constrain b to be positive. This leads to a trading strategy that buys undervalued currencies and sells overvalued currencies relative to purchasing power parity. We will specify and explain the constraint for each model below. For the order flow regressions, we impose no constraints.¹⁰

3.1 Models Conditioning on Order Flow

We estimate three types of predictive regressions that condition on customer order flow. The first type conditions on the order flow of each of the four customer types separately, i.e., asset managers (x_t^{AM}), hedge funds (x_t^{HF}), corporates (x_t^{CO}), and private clients (x_t^{PC}), thus leading to four regressions. The second type conditions on the sum of the four order flows, that is on the aggregate (or total) order flow (TOT): $x_t^{TOT} = x_t^{AM} + x_t^{HF} + x_t^{CO} + x_t^{PC}$. Finally, we condition on all four disaggregated customer order flows in one regression (ALL): $x_t^{ALL} = \{x_t^{AM}, x_t^{HF}, x_t^{CO}, x_t^{PC}\}$. These regressions

¹⁰The constraints fix the sign of the slope of the predictive regression to be the same as the sign expected by theory. This is similar but not identical to Campbell and Thompson (2008) who also impose constraints on the sign of the slope of predictive regressions in assessing predictability in stock returns. In particular, we do not follow Campbell and Thompson (2008) in setting the slope to be equal to zero whenever it has the “wrong” sign as this would effectively mix the random walk model with each predictive regression.

will determine whether there is predictive information in customer order flow and the extent to which different customer groups convey different information.

3.2 Random Walk

We now turn to the set of predictive regressions based on publicly available macroeconomic information. The first specification is the driftless (or naive) random walk (RW) model that sets $a = b = 0$, which is the standard benchmark in this literature since Meese and Rogoff (1983). The RW model forms the basis of the widely used carry trade strategy in active currency management (e.g., Burnside, Eichenbaum and Rebelo, 2009; Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011; Lustig, Roussanov and Verdelhan, 2011; Menkhoff, Sarno, Schmeling and Schrimpf, 2012a).

Note that if the RW holds, then spot exchange rate returns cannot be predicted by macroeconomic information. In contrast, however, FX excess returns are perfectly predictable and equal to the interest rate differential. The portfolio returns of a dynamic FX strategy depend on excess returns not just on spot returns. Hence, even if the RW holds, macroeconomic information is important in determining the returns to a dynamic asset allocation strategy for currencies. We provide more details in the next section.

3.3 Forward Premium and Uncovered Interest Parity

The second specification based on public information is the uncovered interest rate parity (UIP) regression and the third specification is the forward premium (FP) regression. The two specifications are closely related since they share the same predictor:

$$x_t = f_t - s_t, \tag{2}$$

where f_t is the log of the one-period forward exchange rate at time t , which is the rate agreed at time t for an exchange of currencies at $t + 1$. The only difference in the two specifications is that the UIP regression imposes a positive sign on the slope of x_t , whereas the FP regression imposes a negative sign.

The explanation for constraining the slopes with opposite signs in the two specifications is as follows. Under risk neutrality and rational expectations, the UIP condition implies that $a = 0$, $b = 1$, and the error term is serially uncorrelated. However, empirical studies consistently reject this condition and it is a stylized fact that estimates of b often display a negative sign (e.g., Fama, 1984). This implies that the forward rate is a biased estimator of the future spot rate or, alternatively, that high-interest rate currencies tend to appreciate rather than depreciate over time.¹¹

¹¹We implicitly assume that covered interest parity (CIP) holds, so that the interest rate differential is equal to the forward premium, $f_t - s_t = i_t - i_t^*$, where i_t and i_t^* are the domestic and foreign nominal interest rates, respectively. There is ample empirical evidence that CIP holds in practice for the data frequency examined in this paper. For

In this context, the UIP regression constrains b to be positive in order to be qualitatively similar to the UIP condition. Note that we do not constrain b to be equal to one, as would be required for the UIP regression to be fully consistent with the UIP condition. This is because throughout our analysis we only impose a constraint on the sign, not the value, of b . Constraining b to be positive implies that high-interest currencies tend to depreciate on average but not necessarily by an amount that fully offsets the interest rate differential. In contrast, the FP regression exploits a more severe version of the forward bias in exchange rates by constraining b to be negative. This implies that high-interest currencies tend to appreciate on average, consistent with the seminal finding of Fama (1984).

3.4 Purchasing Power Parity

The fourth regression is based on the purchasing power parity (PPP) condition. PPP is a long-run condition, which states that national price levels should be equal when expressed in a common currency (e.g., Taylor and Taylor, 2004). The PPP regression sets x_t as follows:

$$x_t = p_t - p_t^* - s_t, \quad (3)$$

where p_t is the log of the domestic price level and p_t^* the log of the foreign price level. In this regression, we impose a positive sign on b . This is equivalent to a trading strategy that buys undervalued currencies and sells overvalued currencies relative to PPP.

3.5 Monetary Fundamentals

The fifth regression conditions on monetary fundamentals (MF) by specifying x_t as follows:

$$x_t = (m_t - m_t^*) - (y_t - y_t^*) - s_t, \quad (4)$$

where m_t is the log of the domestic money supply and y_t is the log of the domestic real output. Similarly, m_t^* is the log of the foreign money supply and y_t^* is the log of the foreign real output. In this regression, we impose a positive sign on b .

The relation between exchange rates and fundamentals defined in Equation (4) suggests that a deviation of the nominal exchange rate from its long-run equilibrium level determined by current monetary fundamentals requires the exchange rate to move in the future so as to converge towards its long-run equilibrium. The empirical evidence on the relation between exchange rates and fundamentals is mixed. On the one hand, short-run exchange rate variability appears to be disconnected from the underlying monetary fundamentals (e.g., Engel, Mark and West, 2008). On the other hand,

recent evidence, see Akram, Rime and Sarno (2008). The only exception in our sample is the period following Lehman Brothers' bankruptcy, when the CIP violation persisted for a few months (e.g., Mancini-Griffoli and Ranaldo, 2011).

there is growing evidence that exchange rates and monetary fundamentals move together in the long run, which implies a positive value for b in the predictive regression (e.g., Groen, 2000; Rapach and Wohar, 2002).

3.6 Taylor Rule

The sixth specification uses the Taylor (1993) rule (TR). The Taylor rule postulates that a country's central bank raises the short-term nominal interest rate when output is above potential output and/or inflation rises above its desired level. The TR regression defines x_t as follows:

$$x_t = 1.5 (\pi_t - \pi_t^*) + 0.1 (\bar{y}_t - \bar{y}_t^*) + 0.1 (s_t + p_t^* - p_t), \quad (5)$$

where π_t is the domestic inflation rate and $\bar{y}_t$ is the domestic output gap. Similarly, π_t^* is the foreign inflation rate and $\bar{y}_t^*$ the foreign output gap. The output gap is measured as the percent deviation of real output from an estimate of its potential level computed using the Hodrick and Prescott (1997) filter.¹² The parameters on the inflation difference (1.5), output gap difference (0.1) and the real exchange rate (0.1) are fairly standard in the literature (e.g., Engel, Mark and West, 2008; Mark, 2009; Molodtsova and Papell, 2009). In this regression, we impose a positive sign on b .¹³

3.7 Cyclical External Imbalances

The seventh model conditions on a bilateral measure of cyclical external imbalances (NXA) between the US and the foreign country denoted as:

$$x_t = nxa_t. \quad (6)$$

Following Gourinchas and Rey (2007), nxa_t is a global measure of cyclical external imbalances that linearly combines detrended (log) exports, imports, foreign assets and liabilities relative to GDP. In particular, nxa_t is constructed using a two-stage least squares estimator as in Della Corte, Sarno and Sestieri (2012). We first regress the global nxa_t for the US on a constant term and the global nxa_t for the foreign country, and then use the fitted value from this contemporaneous regression as x_t , representing the proxy for the bilateral measure of cyclical external imbalances between the US and the foreign country.

The NXA predictive regression constrains b to be negative, which is justified as follows: a country running external imbalances ($nxa_t < 0$) will experience a future currency depreciation ($\Delta s_{t+1} > 0$)

¹²We construct the output gap using the Hodrick and Prescott (1997) filter using a smoothing parameter equal to 14,400 as in Molodtsova and Papell (2009). For the first period, the output gap is computed using real output data from January 1990 to January 2001. In estimating the Hodrick-Prescott trend out of sample, at any given period t , we only use data up to period $t - 1$. We then update the trend every time a new observation is added to the sample. This captures as closely as possible the information available at the time a forecast is made and avoids look-ahead bias.

¹³We also estimate (rather than fix) the parameters on the inflation difference, output gap difference and the real exchange rate but we find that the results remain qualitatively identical. Hence we use the fixed parameters as above.

that contributes to the process of international financial adjustment through future current account surpluses and/or future higher returns on the net foreign asset portfolio (see Gourinchas and Rey, 2007).

3.8 Momentum

The eighth and final specification captures momentum (MOM) by using the one-month rolling exchange rate return as the conditional mean of the one-period ahead exchange rate. In this regression, we impose a positive sign on b . The MOM strategy produces a long exposure to the currencies that are trending higher, and a short exposure to the currencies that are trending lower. Strictly speaking, momentum is not a strategy based on macroeconomic information. We include it in the set of predictive regressions because, like the carry trade, momentum is a profitable and increasingly popular FX strategy (see Menkhoff, Sarno, Schmeling and Schrimpf, 2012b).

4 Assessing the Predictive Ability of Order Flow

This section describes the framework for evaluating the ability of order flow to predict exchange rate returns in the context of dynamic asset allocation strategies.

4.1 The Dynamic FX Strategy

We design an international asset allocation strategy that involves trading the US dollar vis-à-vis nine major currencies: the Australian dollar, Canadian dollar, Swiss franc, euro, British pound, Japanese yen, Norwegian krone, New Zealand dollar and Swedish krona. Consider a US investor who builds a portfolio by allocating her wealth between ten bonds: one domestic (US), and nine foreign bonds (Australia, Canada, Switzerland, Germany representing the euro area, UK, Japan, Norway, New Zealand and Sweden). The yield of the bonds is proxied by Eurodeposit rates. At each period $t + 1$, the foreign bonds yield a riskless return in local currency but a risky return r_{t+1} in US dollars. The expected US dollar return of investing in a foreign bond is equal to $r_{t+1|t} = i_t^* + \Delta s_{t+1|t}$, where $r_{t+1|t} = E_t[r_{t+1}]$ is the conditional expectation of r_{t+1} , i_t^* is the foreign nominal interest rate, and $\Delta s_{t+1|t} = E_t[\Delta s_{t+1}]$ is the conditional expectation of Δs_{t+1} . Hence the only risk the US investor is exposed to is FX risk.

Each period the investor takes two steps. First, she uses the predictive regressions conditioning on order flow or macroeconomic information to forecast the one-period ahead exchange rate returns. Second, conditional on the forecasts of each model, she dynamically rebalances her portfolio by computing the new optimal weights using the method discussed below. This setup is designed to assess the predictive ability of customer order flow by informing us whether conditioning on order

flow leads to a better performing allocation strategy than conditioning on the random walk or other macroeconomic models.

4.2 Mean-Variance Dynamic Asset Allocation with Transaction Costs

Mean-variance analysis is a natural framework for assessing the economic value of strategies that exploit predictability in the mean and variance. Consider an investor who has a one-period horizon and constructs a dynamically rebalanced portfolio. Computing the time-varying weights of this portfolio requires one-step ahead forecasts of the conditional mean and the conditional variance-covariance matrix. Let r_{t+1} denote the $K \times 1$ vector of risky asset returns at time $t + 1$, $V_{t+1|t} = E_t[(r_{t+1} - r_{t+1|t})(r_{t+1} - r_{t+1|t})']$ the $K \times K$ conditional variance-covariance matrix of r_{t+1} , τ_{t+1} the $K \times 1$ vector of proportional transaction costs, and $\tau_{t+1|t} = E_t[\tau_{t+1}]$ the conditional expectation of τ_{t+1} .

Our analysis focuses on the maximum expected return strategy, which leads to an allocation on the efficient frontier. This strategy maximizes the expected portfolio return at each period t for a given target portfolio volatility:

$$\begin{aligned} \max_{w_t} \quad & r_{p,t+1|t} = w_t' r_{t+1|t} + (1 - w_t' \iota) r_f - \phi_{t+1|t} \\ \text{s.t.} \quad & \sigma_p^* = (w_t' V_{t+1|t} w_t)^{1/2}, \end{aligned} \tag{7}$$

where $r_{p,t+1}$ is the portfolio return at time $t + 1$, $r_{p,t+1|t} = E_t[r_{p,t+1}]$ is the conditional expectation of $r_{p,t+1}$, ι is a $K \times 1$ vector of ones, r_f is the riskless rate, σ_p^* is the target conditional volatility of portfolio returns, and $\phi_{t+1|t}$ is the conditional expectation of the total transaction cost for the portfolio in each period defined as:

$$\phi_{t+1|t} = \sum_{i=1}^K \tau_{i,t+1|t} |w_{i,t} - w_{i,t}^-|, \tag{8}$$

where $w_{i,t}^- = w_{i,t-1} (1 + r_{i,t}) / (1 + r_{p,t})$.

The proportional transaction cost $\tau_{i,t+1|t}$ for each asset i is computed as follows. We first define the excess return of holding foreign currency for one period net of transaction costs as:

$$er_{t+1}^{net} = s_{t+1}^b - f_t^a, \tag{9}$$

where s_{t+1}^b is the bid-quote for the spot rate at time $t + 1$, and f_t^a is the ask-quote for the forward rate at time t . This is the excess return for an investor who buys a forward contract at time t for exchanging the domestic currency into foreign currency at time $t + 1$, and then, at time $t + 1$ she converts the proceeds of the forward contract back into the domestic currency at the $t + 1$ spot rate.

We can rewrite the above expression using mid-quotes to obtain:

$$\begin{aligned} er_{t+1}^{net} &= \left(s_{t+1} - \frac{s_{t+1}^a - s_{t+1}^b}{2} \right) - \left(f_t + \frac{f_t^a - f_t^b}{2} \right) \\ &= (s_{t+1} - f_t) - c_{t+1}, \end{aligned} \quad (10)$$

where s_{t+1} and f_t are the mid-quotes for the spot and forward exchange rate, and $c_{t+1} = (s_{t+1}^a - s_{t+1}^b + f_t^a - f_t^b)/2$ represents the round-trip proportional transaction cost of the simple trading strategy. In our setup, we define $\tau_{t+1} = c_{t+1}/2$ as the one-way proportional transaction cost for increasing or decreasing the portfolio weight at time $t + 1$ on a given foreign currency.

In the empirical implementation of the mean-variance strategy, we compute the time-varying weights w_t using information up to time t . These weights determine the $t + 1$ portfolio return $r_{p,t+1}$. However, the transaction cost τ_{t+1} relevant to $t + 1$ returns will only be known ex post, whereas the weights (which require an estimate of τ_{t+1}) are set ex ante. We avoid this complication by estimating $\tau_{t+1|t}$ using the 3-month rolling average of the times series of τ_t using information up to time t . Note that to compute the weight w_t we use the estimate $\tau_{t+1|t}$, but to compute the net portfolio returns, which are known ex post, we use the realized τ_{t+1} value.

The inclusion of transaction costs in the mean-variance optimization implies that the solution for the time-varying weights w_t is not available in closed form but is obtained via numerical optimization.¹⁴ Once the optimal weights are computed, the return on the investor's portfolio net of the realized transaction costs is equal to:

$$r_{p,t+1} = w_t' r_{t+1} + (1 - w_t' \iota) r_f - \phi_{t+1}. \quad (11)$$

Finally, note that we do not model the dynamics of FX return volatility and correlation. For the in-sample analysis, we set the conditional covariance matrix to be equal to the sample covariance matrix using the full data sample. For the out-of-sample analysis, we do the same using data up to the time the forecast is made in order to avoid look-ahead bias. This approach implies that the optimal weights will vary across the empirical exchange rate models only to the extent that the predictive regressions produce different forecasts of the exchange rate returns.¹⁵

¹⁴We use a linear transaction cost function because it can be solved globally and efficiently as a convex portfolio optimization problem. In practice, transaction costs may be a concave function of the amount traded. This happens, for example, when there is an additional fixed component to allow the total transaction cost to decrease as the amount traded increases. This portfolio optimization problem cannot be solved directly via convex optimization (see Lobo, Fazel and Boyd, 2007).

¹⁵See Della Corte, Sarno and Tsiakas (2009, 2012) for an empirical analysis of the effect of dynamic volatility and correlation on FX trading strategies. See also Berger, Chaboud, Hjalmarsson and Howorka (2009) for the relation between FX volatility and order flow.

4.3 Performance Measures

We evaluate the performance of the exchange rate models using the Goetzmann, Ingersoll, Spiegel and Welch (2007) performance measure defined as:

$$M(r_p) = \frac{1}{(1-\gamma)} \ln \left\{ \frac{1}{T} \sum_{t=1}^T \left(\frac{1+r_{p,t}}{1+r_f} \right)^{1-\gamma} \right\}, \quad (12)$$

where $M(r_p)$ is an estimate of the portfolio's premium return after adjusting for risk, which can be interpreted as the certainty equivalent of the excess portfolio returns. This is an attractive criterion since it is robust to the distribution of portfolio returns and does not require the assumption of a particular utility function to rank portfolios. The parameter γ denotes the investor's degree of relative risk aversion (RRA).

We compare the performance of the exchange rate model conditioning on order flow to the benchmark RW by computing the difference:

$$\mathcal{P} = M(r_p^*) - M(r_p^b), \quad (13)$$

where r_p^* are the portfolio returns of the order flow strategy and r_p^b are the portfolio returns of the benchmark RW. We interpret $\mathcal{P}$ as the maximum performance fee an investor will pay to switch from the RW to the order flow strategy. In other words, this performance criterion measures how much a mean-variance investor is willing to pay for conditioning on better exchange rate forecasts. We report $\mathcal{P}$ in annualized basis points (bps).¹⁶

In the context of mean-variance analysis, perhaps the most commonly used performance measure is the Sharpe ratio ($\mathcal{SR}$). The realized $\mathcal{SR}$ is equal to the average excess return of a portfolio divided by the standard deviation of the portfolio returns. We also compute the Sortino ratio ($\mathcal{SO}$), which measures the excess return to “bad” volatility. Unlike the $\mathcal{SR}$, the $\mathcal{SO}$ differentiates between volatility due to “up” and “down” movements in portfolio returns. It is equal to the average excess return divided by the standard deviation of the negative returns. In other words, the $\mathcal{SO}$ does not take into account positive returns in computing volatility because these are desirable. In addition, we report the maximum drawdown ($\mathcal{MDD}$), which is the maximum cumulative loss from the strategy's peak to the following trough. As large drawdowns usually lead to fund redemptions, it follows that a reasonably low $\mathcal{MDD}$ is critical to the success of any asset allocation strategy.

¹⁶Note that the premium return $\mathcal{P}$ defined above gives very similar results to the Fleming, Kirby and Ost diek (2001) performance fee based on quadratic utility that is often used in the literature.

5 Empirical Results

5.1 The Correlation between Currency Order Flow and Exchange Rate Returns

We begin our empirical analysis by calculating the correlation between contemporaneous currency order flow and exchange rate returns at different horizons, as in Froot and Ramadorai (2005). This simple measure provides a preliminary way of assessing the statistical significance of the relation between order flow and FX returns at different horizons. For example, a significant positive correlation at the daily horizon suggests that daily order flow can explain daily FX returns. More interestingly, a significant positive correlation at a longer horizon indicates that flows must either positively anticipate returns or positively follow recent returns, which is therefore evidence of predictability in at least one direction.

Figure 1 displays for each customer type the average correlation across all exchange rates. Using the average correlation across nine exchange rates significantly increases statistical power, which is useful for long horizons as fewer observations are available. The horizon is reported in log-scale on the horizontal axis, running from the 1-day horizon (10^0 days) to 252 days ($> 10^2$). To assess the significance of these correlations, the figure also shows the 90% confidence intervals generated by 10,000 bootstrap replications under the null hypothesis that each order flow and FX return series is independent and identically distributed (i.i.d.).¹⁷

Consistent with the daily cross-correlations reported in Table 3, the correlations at the 1-day horizon are positive for AM and HF flows, and negative for CO and PC flows. For AM flows, the positive correlations increase markedly with horizon and are highly statistically significant as they far exceed the bootstrap confidence intervals. For HF flows, the correlations are positive but are only significant for horizons of one to two weeks. For both AM and HF, the correlations eventually decrease, which may reflect reversals in their positions. Note, however, that despite decreasing, the AM and HF correlations remain positive, and for AM in particular the decay starts at very long horizons. For CO flows, the initial 1-day correlations are negative, then decrease further and remain significant for up to a month, and then turn positive and insignificant. Finally, the PC correlations are negative and highly significant for all medium horizons up to two months.

Overall, this preliminary analysis provides evidence that for short and medium horizons of up to a few weeks the order flow of asset managers (and to a lesser extent hedge funds) may have high positive predictive ability, whereas the order flow of corporate and especially of private clients may have negative predictive ability. These results are tentative since we formally assess the predictive ability of customer order flow later. Still, the findings are certainly consistent with the notion that

¹⁷Note that in contrast to Froot and Ramadorai (2005) who use a Monte Carlo procedure to derive the confidence bounds, we use bootstrapping to avoid the assumption of a normal distribution.

asset managers and hedge funds are informed investors and their positive order flow tends to be associated with positive future exchange rate returns. In contrast, corporates and especially private clients are less informed investors, and possibly liquidity providers, whose positive order flow tends to be associated with negative future exchange rate returns. Next we turn to exploring whether the information content in order flow has predictive power for future exchange rate returns in the context of dynamic asset allocation.

5.2 The Predictive Ability of Order Flow

This section discusses the empirical results relating to the predictive ability of customer order flow. Our approach is based on a dynamic asset allocation strategy that conditions on the four types of customer order flow. It is natural to use a trading strategy to assess the predictive ability of customer order flow since it is trades that reflect the customers' information. We evaluate each strategy in terms of the returns it generates. This allows us to measure the tangible economic gains and hence the value of conditioning on different types of customer order flow. It is important to recall that customer transactions are private and the details of each trade are known in real time only to the dealer customers transact with. Therefore, if customer order flows have predictive information for future FX excess returns, this information is not widely available to market participants. With this in mind, our main objective is to use a trading strategy as the tool for determining whether predictive information is conveyed to dealers by customer order flows, not as a recommended method for profitable FX trading.¹⁸

We analyze the performance of dynamically rebalanced portfolios based on the order flow strategies relative to the random walk (RW) benchmark. We use monthly rebalancing to be consistent with the analysis in the next section that utilizes monthly macroeconomic variables. In this setting, the investor obtains forecasts of exchange rate returns for next month conditioning on order flow information available at the time of the forecast; she then chooses investment weights using the maximum expected return strategy for an annual target volatility of $\sigma_p^* = 10\%$ and a coefficient of relative risk aversion $\gamma = 6$.¹⁹

Forecasting and portfolio optimization are conducted both in sample and out of sample. The in-sample prediction uses the predictive regressions described in Section 3 estimated over the full data sample ranging from January 2001 to May 2011. For the out-of-sample analysis, we first estimate the predictive regressions over the initial period of January 2001 to December 2003, and

¹⁸Note, however, that the order flow information used in this paper is obviously known in real time to UBS, which can use this information legitimately provided it is done according to standard compliance rules.

¹⁹The choice of σ_p^* and γ is reasonable and consistent with numerous previous empirical studies. We have experimented with different σ_p^* and γ values and found that qualitatively they have little effect on the asset allocation results. Note that for the maximum return strategy, γ only affects the performance evaluation (in the calculation of the premium return) but not the asset allocation.

then reestimate these regressions recursively until the end of the full sample (May 2011). Each out-of-sample prediction is conditional on information available at the time of the forecast.

Our assessment of the performance of the models focuses on the realized excess portfolio returns and their descriptive statistics, the Sharpe ratio ($\mathcal{SR}$), the Sortino ratio ($\mathcal{SO}$), the maximum draw-down ($\mathcal{MDD}$) and the performance fee ($\mathcal{P}$). In evaluating the profitability of dynamic strategies, the effect of transaction costs is an essential consideration. For instance, if the bid-ask spread in trading currencies is sufficiently high, the order flow strategies may be too costly to implement. Hence all asset allocation results are reported net of transaction costs, where the bid-ask spread is explicitly taken into account in the optimization that delivers the mean-variance weights. This is an important aspect of our analysis as the transaction costs endogenously determine the optimal asset allocation. We consider an effective transaction cost that is equal to 50% of the quoted spread.²⁰

We now turn to the evaluation of the predictive ability of the models. The results for monthly rebalancing are in Table 4. Focusing on the out-of-sample results, two findings are noteworthy. First, the order flow models tend to significantly outperform the RW benchmark thus indicating that order flow has substantial predictive information at monthly horizon. The best performing model is AM, which has a Sharpe ratio net of transaction costs of 0.84 (compared to 0.21 for RW) and a performance fee of 879 annual basis points (bps). The HF model also performs well with $\mathcal{SR} = 0.61$ and $\mathcal{P} = 560$ bps. CO follows in third place with $\mathcal{SR} = 0.51$ and $\mathcal{P} = 345$ bps. PC is the only type of order flow that produces a negative Sharpe ratio and a negative performance fee. This establishes that it is important to account for the heterogeneity in customer order flow as the AM and HM order flows, in particular, exhibit substantial out-of-sample predictive ability.

Second, the performance of the disaggregated predictive regression (ALL) is far superior to that of the total (aggregate) order flow model (TOT). For example, out of sample TOT produces $\mathcal{SR} = 0.40$ and $\mathcal{P} = 190$ bps, whereas ALL delivers $\mathcal{SR} = 0.79$ and $\mathcal{P} = 801$ bps. In other words, by conditioning on all individual customer order flows separately, we can generate substantial economic gains. In conclusion, these results show that a risk-averse investor would be willing to pay a large performance fee to switch from a random walk strategy to a strategy that conditions on order flow information as the out-of-sample performance fee can be as high as 879 basis points per year. This is strong evidence of the predictive power of order flow information as compared to the RW benchmark.

An explanation for the very good performance of the AM monthly order flow could be, for example, that asset managers are large investors with long horizons. It is indeed likely that informed traders choose to trade gradually, hiding their trades among the uninformed trades, in order to avoid

²⁰It is well-documented that the effective spread is generally lower than the quoted spread, since trading will take place at the best price quoted at any point in time, suggesting that the worse quotes will not attract trades (e.g., Lyons, 2001; Della Corte, Sarno and Tsiakas, 2011).

a rapid adjustment in the price they pay.²¹ In practice, AM may hold large positions which take time to unwind. Hence the monthly frequency is well-suited in capturing the predictive information of AM order flow. On the other extreme, the evidence suggests that PC trade often and in the wrong direction.²² A possible explanation for this result is that PC are thought of as liquidity providers, who are rather unsophisticated relative to AM and HF.

To provide a visual illustration of the results, Figure 2 plots the out-of-sample cumulative wealth for the four individual order flow models relative to the RW benchmark starting with an initial wealth of \$1. Consistent with the results above, the figure shows that the cumulative wealth for the AM, HF and CO order flow strategies is higher than the RW, but for PC it is lower. More importantly, the figure illustrates that while at the beginning all order flow models tend to comove with the RW, after the crisis in 2007 the AM and HF models, in particular, outperform considerably the RW. In contrast, from the end of 2008 onwards, when Lehman Brothers collapsed, PC significantly underperforms relative to the RW. Hence it would be reasonable to conclude that much of the order flow prior to the crisis is driven by carry positions, whereas after the crisis it is not. This is not surprising as the unwinding of carry trades during the crisis would have reduced the role of the carry trade in determining order flow. Furthermore, the performance of each customer group seems to depend on how it reacted to the crisis.

5.3 Order Flow and Macroeconomic Information

5.3.1 The Empirical Framework

Having established the predictive ability of customer order flow, next we examine the extent to which this predictive information is related to widely available macroeconomic information. In other words, we ask whether macroeconomic information drives customer order flow for currencies. We address this question by setting up a framework based on trading strategies that condition on order flow and on the macroeconomic models. In particular, we investigate whether the excess portfolio returns generated by the strategies that condition on order flow are correlated with the excess portfolio returns generated by eight alternative strategies: random walk (RW), forward premium (FP), uncovered interest parity (UIP), purchasing power parity (PPP), monetary fundamentals (MF), Taylor rule (TR), cyclical external imbalances (NXA) and momentum (MOM). Note that, under the random walk, macroeconomic information determines the FX excess returns but not the FX spot returns. Furthermore, we include momentum not because it is based on macroeconomic information but because, like the carry trade, it is a profitable and increasingly popular FX strategy. Finally, recall

²¹This is consistent with the multiple-period version of the Kyle (1985) model. See, for example, Chapter 4 in Lyons (2001) for a more detailed description of this idea.

²²PC do trade more often since in unreported results we find that PC portfolios have a higher turnover than AM and HF portfolios.

that all macroeconomic variables are constructed using real-time data that was available to market participants at the time forecasts are made.

Our empirical approach involves three steps. First, we estimate a set of predictive regressions conditioning on order flow and macroeconomic models that deliver out-of-sample one-month ahead forecasts for exchange rate returns. We focus on monthly returns as this is the frequency at which most macroeconomic information is released. The out-of-sample forecasts are for the period of January 2004 to May 2011. As mentioned earlier, in this step we constrain the sign of the slope parameters in the predictive regressions to be consistent with the economic theory that each of the regressions represents: the slope parameters are set to be positive for UIP, TR, PPP, MF and MOM, and negative for FP and NXA. This guarantees that we capture the intended macroeconomic effects and is similar to the constraints imposed by Campbell and Thompson (2008) in assessing predictability in stock returns. Second, we use these forecasts in the mean-variance dynamic asset allocation to generate the excess portfolios net of transaction costs. The bid-ask spreads endogenously affect the weights in all strategies. Third, we regress the excess portfolio returns of the order flow strategy for each customer type on the excess portfolio returns of the eight alternative strategies.

This empirical framework uses the realized portfolio returns of strategies that condition on the expected returns (i.e., forecasts) generated by order flow and macroeconomic information. Alternatively, one could focus on the expected returns of the order flow and macroeconomic strategies. We choose to focus on realized (as opposed to expected) returns because this is both a more natural and a more stringent test of whether there is a link between order flow and macroeconomic information. To elaborate, in a predictive framework investors typically make trading decisions using expected returns but evaluate these decisions using realized returns. As a result of using such a framework, not only do we relate order flow to macroeconomic information but we do so for the specific information subsets of these strategies that have predictive ability, and hence value, to an investor. Without question, to relate order flow to macroeconomic information in a predictive framework that uses realized returns sets a higher bar than using expected returns. This is simply because expected returns on their own do not indicate whether the information used to form expectations is actually valuable. For these reasons, in what follows we adopt a predictive empirical framework based on realized portfolio returns.

In addition to helping us understand what drives order flow, this framework sheds light on questions such as: what strategies do asset managers, hedge funds, corporates and private clients follow? Are these strategies different among customer types? What type of macroeconomic information drives the currency demand and supply (captured by order flow) of different customer groups? Can the predictive information content of order flow be explained solely using macroeconomic information, or does it contain additional information that cannot be recovered with a combination of public

information?

These questions are central to two lines of research. First, we can understand better the information reflected in the decisions of customers on what currencies to buy and sell over time. This is therefore related to the broad literature on the behavior of FX currency managers, their performance and risk exposure (e.g., Pojarliev and Levich, 2008, 2010, 2012). The main difference to our study is that this literature tends to focus on directly observed returns of (say) asset managers and hedge funds, in an attempt to replicate these returns and assess whether they provide an “alpha” due to skill or superior information. In contrast, our study conditions on trading decisions based on customer order flow.

Second, recent theoretical literature formalizes the notion that order flow conveys fundamental information about exchange rates by aggregating trading decisions based on economic information dispersed among the agents in the economy (e.g., Evans and Lyons, 2008; Bacchetta and van Wincoop, 2006). This implies that order flow ought to be empirically related to macroeconomic information, and the former effectively summarizes the latter.

5.3.2 The Empirical Results

Table 5 reports the regression results for the monthly excess portfolio returns of each customer flow on the monthly excess portfolio returns of the eight macroeconomic information strategies. The out-of-sample period of January 2004 to May 2011 comprises 89 monthly observations. For each customer type, we estimate four regressions: one where all eight macroeconomic strategies are used and three where only one of RW, UIP and TR are used one at a time together with the remaining ones. This is because the returns of the RW, UIP and TR strategies are highly correlated. We compute bootstrapped standard errors and p -values obtained by resampling 10,000 times the portfolio weights by means of moving block bootstrap (Gonçalves and White, 2005).

We find that the excess returns generated from the order flow strategies can to a large extent be explained by a combination of the eight macroeconomic strategies. The main results can be summarized as follows. For the regression that includes all eight macroeconomic strategies, the $\bar{R}^2$ ranges from 17.9% (AM) to 48.4% (CO). The $\bar{R}^2$ for AM and HF is lower than for CO and PC indicating a lower dependence of the AM and HF trading strategies on macroeconomic information. Furthermore, the $\bar{R}^2$ consistently increases when we include the RW model in the regressions, which indicates that interest rates are a crucial determinant of the predictive ability of order flow. Note that, with few exceptions, the betas on the macroeconomic strategies tend to be positive but insignificant, an issue we investigate further below. The alphas are positive for AM, HF and CO, and negative for PC, but in all cases they are insignificant.

Together these results imply that trading strategies conditioning on customer order flow can be

explained to a large extent by particular combinations of macroeconomic information. The absence of a significant (positive) alpha further suggests that there is no additional information in customer order flow that is unrelated to macroeconomic news. We conclude, therefore, that order flow is related macroeconomic information but offers no additional information over and above widely available public information. This is a new and important result in this literature that further justifies the use of order flow as the conduit through which macroeconomic information is transmitted to exchange rates.²³

Finally, it is important to note that these results are only known *ex post*. We need to first observe order flow and the macroeconomic variables in real time, and then we can explain which type of macroeconomic information drives order flow and with what weight. Only once the time series of out-of-sample portfolio returns are known we can estimate the relation between the order flow strategies and the macroeconomic strategies. Furthermore, the betas in these regressions tend to be insignificant. This, for example, could be due to time-variation and structural breaks in the relation of order flow to fundamentals. In what follows, we explore this issue in detail.

5.3.3 Time-Varying Parameter Regressions

While the regressions capturing the relation between order flow and macroeconomic information are estimated with constant parameters, it is likely that these parameters change over time. This, for example, could be due to structural breaks in the relation of order flow to fundamentals due to the recent crisis. It is also well documented that FX participants change over time the weight they assign to different fundamentals. This is consistent with the scapegoat theory of Bacchetta and van Wincoop (2004, 2011), where over time the market may focus its attention on a different macroeconomic variable (the scapegoat). This happens when traders assign a different weight to a macroeconomic indicator over time as the market rationally searches for an explanation for the observed exchange rate change.

We investigate this possibility in more detail by estimating time-varying parameter (TVP) regressions using the same inputs as before: regress monthly portfolio excess returns generated by conditioning out-of-sample on order flows (left-hand side) on monthly portfolio excess returns generated by conditioning out-of-sample on the eight macroeconomic strategies (right-hand side). Note that we impose the same constraints on the sign of the slopes of the predictive regressions as before. The TVP regressions are estimated with Bayesian methods that are well suited to small samples, as

²³It is interesting to note that our empirical results are consistent with the results of Pojarliev and Levich (2008, 2010, 2012), who find that the majority of the successful currency managers are “beta grazers” rather than “alpha hunters” as they implement strategies with high betas on risk factors rather than high alpha strategies. The alphas of most currency managers are generally low and insignificant. For the few cases that alphas are high and significant, they are not persistent as good performance in one year is not correlated to good performance in the next year.

described in the Appendix.

We begin with Figure 3, which plots the time-varying alphas over time. We have seen previously that assuming constant alphas leads to positive (average) estimates for AM, HF and CO, negative for PC and in all cases they are insignificant. With time-varying alphas a more complete picture emerges as we can identify when alphas are positive, when they are negative as well as when they are significant. For example, after the collapse of Lehman Brothers in September 2008, the AM and HF order flow strategies delivered a significant positive alpha, thus performing better than a combination of macroeconomic fundamentals. In contrast, over the same period the PC strategy delivered a significant negative alpha. Clearly, therefore, more sophisticated investors seem to have significantly outperformed the less sophisticated investors during the crisis.

The next four figures (Figs. 4 to 7) display the time-varying betas and time-varying R_t^2 for each customer group.²⁴ There are several interesting results. The betas vary considerably over time thus justifying the use of TVP regressions. During the crisis, for example, asset managers moved out of the random walk (i.e., carry trade) and into the Taylor rule and purchasing power parity. Private clients also moved out of the random walk as well as out of the Taylor rule and into purchasing power parity. In essence, these graphs show the anatomy of the month-by-month trading strategies of the four customer groups. Note also that the R_t^2 tends to be high most of the time for all customer groups, typically in the range of 80%, but exhibits negative spikes. Indeed, the average R_t^2 is higher in the TVP regressions than in the constant parameter regressions: 50% for AM, 60% for HF, and 70% for CO and PC, compared to 18%, 22%, 48% and 44%, respectively, in the constant parameter regressions. This is further evidence of the strong relation between order flow and macroeconomic information. Note that in these TVP regressions, the main point is to identify the periods when the alphas may be significant, display the time-variation in the betas and show that the average R^2 is high.

We conclude our empirical analysis with a set of figures that provide another way of looking at the time-varying relation between order flow and macroeconomic information. Figures 8 to 11 plot the cumulative monthly order flow of a given customer group associated with a particular macroeconomic strategy. This cumulative order flow is defined as the cumulative sum over time of the product of the order flow for each currency times the mean-variance weight for that currency implied by a macroeconomic strategy. It indicates the amount of money that (say) asset managers would have invested in a macroeconomic strategy over time and hence captures the exposure to a particular set of macroeconomic fundamentals. The figures plot the cumulative order flow against the cumulative

²⁴We compute a time-varying measure of the coefficient of determination R_t^2 using information from the Kalman filter implemented in the Bayesian estimation algorithm. Conditional on the information of each period t , the time-varying R_t^2 accounts for the variance of the residuals as well as the covariance matrix of the time-varying beta parameters. For more details, see the Appendix.

monthly wealth of the relevant macroeconomic strategy to provide a visual illustration of whether the two are correlated over time. It is clear from the figures that the cumulative order flow fluctuates over time and seems highly correlated with the performance of the macroeconomic strategies. For example, it is noteworthy that asset managers and hedge funds reduced their exposure to the carry trade at least a few months before the crisis started.

5.4 Can We Replicate Order Flow in Real Time?

In this section we ask whether we can replicate in real time the predictive information of order flow for exchange rate returns. This is an important question that we address as follows. We take the point of view of an investor who can optimally combine the exchange rate return forecasts of the eight macroeconomic information strategies using standard forecast combination methods. We then assess whether the investor can replicate *ex ante* the predictive ability of order flow. If so, order flow simply replicates a particular forecast combination and does not contribute any further predictability. If not, the way order flow reflects macroeconomic information can only be revealed *ex ante* by order flow itself. In the latter case, order flow summarizes the available macroeconomic information in a distinct and effective manner, as suggested by microstructure theory.

Following a long line of research since the seminal work of Bates and Granger (1969), we provide a comprehensive analysis of the performance of a number of standard forecast combinations.²⁵ We estimate $N = 8$ predictive regressions for a vector of $K = 9$ exchange rates. Each predictive regression $j \leq N$ generates an individual forecast $\Delta s_{j,t+1|t}$ for the vector of one-step ahead exchange rate returns. We define the combined forecast $\Delta s_{c,t+1|t}$ as the weighted average of the N individual forecasts:

$$\Delta s_{c,t+1|t} = \sum_{j=1}^N \theta_{j,t} \Delta s_{j,t+1|t}, \quad (14)$$

where $\{\theta_{j,t}\}_{j=1}^N$ are the *ex-ante* combining weights determined at time t .

The combining methods we consider differ in how the weights are determined and can be organized into four types. The first type uses simple averaging schemes: mean, median, and trimmed mean. The mean (AVE) combination forecast sets $\theta_{j,t} = 1/N$ in Equation (14); the median (MED) combination forecast is the median of $\{\Delta s_{j,t+1|t}\}_{j=1}^N$; and the trimmed (TRI) mean combination forecast sets $\theta_{j,t} = 0$ for the individual forecasts with the smallest and largest values, and $\theta_{j,t} = 1/(N - 2)$ for the remaining individual forecasts in Equation (14). These combined forecasts disregard the historical performance of the individual forecasts.

The second type of combined forecasts is based on Bates and Granger (1969) and Stock and

²⁵For recent applications of combined forecasts see Stock and Watson (2004), Timmermann (2006) and Rapach, Strauss and Zhou (2010). An alternative way to combine forecasts from different models is to use the method of Barroso and Santa-Clara (2012) for designing optimal currency portfolios.

Watson (2004), and uses statistical information on the past performance of each individual model. In particular, it sets the weights by computing the following forecast combination based on the mean squared error (MSE):

$$\theta_{j,t} = \frac{MSE_{j,t}^{-1}}{\sum_{j=1}^N MSE_{j,t}^{-1}}, \quad MSE_{j,t} = \frac{1}{T} \sum_{t=1}^T (\Delta s_{j,t} - \Delta s_{j,t|t-1})^2. \quad (15)$$

The third type follows the Welch and Goyal (2008) “kitchen sink” (KS) regression that incorporates all N predictive variables $\{x_t^j\}_{j=1}^N$ in one predictive regression:

$$\Delta s_{c,t+1} = a + \sum_{j=1}^N b_j x_t^j + \varepsilon_{t+1}. \quad (16)$$

The fourth and final type implements Principal Component Analysis (PCA) and uses the first principal component.

Our empirical analysis computes the AVE, MED, TRI, MSE, KS and PCA combined forecasts using the eight individual forecasts of the macroeconomic models. We then compare the performance of the forecast combinations to the RW benchmark. The objective of this exercise is to assess whether standard methods of combining macroeconomic information can capture the predictive information of order flow. An important advantage of these forecast combination methods is that they produce ex ante forecasts that can be used in realistic trading strategies.

The results, reported in Table 6, provide overwhelming evidence that none of the forecast combination methods can outperform the RW benchmark out of sample. All forecast combinations display negative Sharpe ratios and highly negative performance fees. This is in stark contrast to the predictive ability of order flow reported in Table 4. Although order flow significantly outperforms the RW out of sample, none of forecast combinations do so. In fact, the best performing combined strategy is MED with a performance fee $\mathcal{P} = -1240$ bps. In contrast, $\mathcal{P} = 801$ bps for the order flow strategy conditioning on the information of all four customer groups (ALL). Hence order flow outperforms the best combined strategy by approximately 2,000 bps or 20% per year. We conclude, therefore, that order flow aggregates the available macroeconomic information in a way that cannot be replicated ex ante by standard forecast combination methods. In short, our findings suggest that the answer to the original question on whether we can replicate in real time the predictive ability of order flow is simply no.

5.5 Summary of Results

The empirical evidence supports the following three main findings. First, our results establish that customer order flow has substantial predictive ability for exchange rate returns. This predictability can lead to profitable trading strategies net of transaction costs. The most valuable information is in the customer order flow of asset managers and hedge funds.

Second, combinations of empirical exchange rate models based on macroeconomic information can explain *ex post* a large part (on average up to 50%) of the predictive ability of order flow. Therefore, macroeconomic information is “price relevant” by generating currency orders and informing customers’ trading decisions. Furthermore, the relation between order flow and macroeconomic information can vary significantly over time as investors react to changes in the macroeconomic environment by assigning a different weight each month to different macroeconomic variables. This result is consistent with the scapegoat theory of Bacchetta and van Wincoop (2004, 2011). In particular, time-varying parameter regressions provide a more complete picture of how customer order flow relates to different macroeconomic fundamentals. In these regressions, macroeconomic information can on average explain 50% to 70% of the predictive ability of order flow. Notably, with the exception of a short period around the crisis, there is no evidence that order flow has additional predictive information over and above the information embedded in macroeconomic fundamentals.

Third, standard forecast combinations of macroeconomic variables fail to replicate *ex ante* the predictive ability of order flow. In fact, order flow strategies can outperform the best forecast combination based on macroeconomic information by approximately 20% per year. This leads us to finally conclude that order flow provides a distinct and effective way of aggregating macroeconomic information.

6 Conclusions

Trades between customers and FX dealers generate currency order flow that conveys the customers’ information on exchange rates. Market microstructure theory suggests that the information in customer order flow is an aggregation of dispersed macroeconomic information. In this paper, we examine empirically whether this is indeed the case. Our analysis is based on a proprietary data set on order flow that covers the G10 currencies for a sample of over ten years. The order flow data represent the transactions of UBS, a leading FX dealer globally, and four groups of customers: asset managers, hedge funds, corporate clients, and private clients. We also manually compile a data set on real-time macroeconomic indicators, which mimics the public information set available to investors over the sample.

Our findings support the view that currency order flow is a very useful summary of widely available public information that has significant predictive ability for exchange rates. Simply put, currency trades by different customers are to a large degree based on a combination of macroeconomic information available in real time. There are three aspects of the relation between the information in order flow and macroeconomic information that are worth noting. First, this relation changes over time, especially during the recent crisis. For example, before the crisis the carry trade was

a prominent driver of customer trades, and after the crisis it is not. Second, with the benefit of hindsight, we do not find evidence that customers have any information over and above what is publicly available. The more successful customers (notably asset managers and hedge funds) simply make better use of macroeconomic information available to everyone. Third, we can only establish a relation between the information in order flow and macroeconomic information ex post, after we observe the customer trades. We cannot establish this relation ex ante. It is very hard, for example, to replicate in real time the information content of the order flow of asset managers and hedge funds simply by combining macroeconomic variables. This is not because these customers have superior information but because they may have a superior interpretation of public information that can only be observed ex post through order flow.

These findings suggest that currency order flow reflects public information but is still key to exchange rate determination because it has significant out-of-sample forecasting ability that we cannot otherwise replicate. In a nutshell, we conclude that order flow provides a distinct and effective way of aggregating dispersed public information about economic fundamentals that are relevant to exchange rates.

Appendix: Bayesian Estimation of the Time-Varying Parameter Model

The regression model with time-varying parameters is specified as follows:

$$y_t = x_t' \beta_t + \varepsilon_t \quad (17)$$

$$\beta_{t+1} = \beta_t + u_t, \quad (18)$$

where y_t is the dependent variable, $x_t = (1, x_{1,t}, \dots, x_{N-1,t})'$ is an $N \times 1$ vector of independent variables, $\beta_t = (\alpha_t, \beta_{1,t}, \dots, \beta_{N-1,t})'$ is an $N \times 1$ vector of unobservable time-varying coefficients, and ε_t and u_t are independent innovations, for $t = 1, \dots, T$. The random error ε_t in the measurement equation (17) is assumed to be normally distributed with zero mean and constant precision h : $\varepsilon_t \sim N(0, h^{-1})$. The vector β_t in the state equation (18) evolves randomly and is exposed to uncorrelated shocks u_t that are jointly normal with zero mean and diagonal precision matrix H : $u_t \sim N(0, H^{-1})$. The innovations ε_t and u_s are independent of one another for all s and t .

We estimate the precision h , the diagonal precision matrix H , and the time-varying coefficients β_t via Bayesian inference by constructing a Markov chain whose limiting distribution is the target posterior density. This Markov chain is a Gibbs sampler in which all parameters are drawn sequentially from their full conditional posterior distribution. The chain is then iterated and the sampled draws, beyond a burn-in period, are treated as variates from the target posterior distribution. We define the joint prior as $p(h, H, B) = p(h) p(H) p(B|H)$, where B is a $T \times N$ matrix that stacks together all time-varying coefficients. For h we assume a Gamma prior $G\left(\frac{\nu}{2}, \frac{2}{\nu s^2}\right)$ with mean $\underline{s}^{-2} = 1$ and degree of freedom $\underline{\nu} = 2$. For H we assume a Wishart prior $W(\underline{\nu}_H, \underline{H})$ with scale matrix $\underline{H} = I_N$ and degree of freedom $\underline{\nu}_H = N$.²⁶ The prior for B is obtained through the state equation (18) as $p(B|H) = p(\beta_1|H) p(\beta_2|\beta_1, H) \dots p(\beta_T|\beta_{T-1}, H)$ where $\beta_{t+1}|\beta_t, H \sim N(\beta_t, H^{-1})$, and $\beta_1|H \sim N(0, H^{-1})$. The Gibbs sampling algorithm is summarized below (for more details see Koop, 2003):

1. Initialize h and H .
2. Sample the matrix B in one block from the conditional posterior distribution $B|y, h, H$ using de Jong and Shephard (1995) algorithm, which is designed for efficient sampling of the state vector in a state space model.
3. Sample $H|y, B \sim W(\bar{\nu}_H, \bar{H})$, where $\bar{H} = \left[\underline{H}^{-1} + \sum_{t=1}^T (\beta_{t+1} - \beta_t) (\beta_{t+1} - \beta_t)' \right]^{-1}$ and $\bar{\nu}_H = T + \underline{\nu}_H$.

²⁶Since the errors are uncorrelated with one another in the state equation, the Wishart prior simplifies to N Gamma priors, that is a gamma prior for each diagonal element of H . See Koop (2003) for more details.

4. Sample $h|y, B \sim G\left(\frac{\bar{\nu}}{2}, \frac{2}{\bar{\nu}\bar{s}^2}\right)$, where $\bar{s}^2 = \left[\sum_{t=1}^T (y_t - x_t'\beta_t)^2 + \underline{\nu}s^2\right] \bar{\nu}^{-1}$ and $\bar{\nu} = T + \underline{\nu}$.
5. Go to step 2 and iterate 100,000 times beyond a burn-in of 20,000 iterations.

The algorithm is a Gibbs sampler with data augmentation that sequentially draws from $p(B|y, h, H)$, $p(H|y, B)$, and $p(h|y, B)$. While the last two conditional posterior distributions are straightforward, the first conditional posterior distribution is less so as the state variables have high dimension and are not independent across time. We use the procedure developed by de Jong and Shephard (1995) which runs the Kalman filter and the simulation smoother. The Kalman filter is a recursive algorithm that calculates the mean and the covariance of the unobserved variables through prediction and updating steps (e.g., Durbin and Koopman, 2012). The simulation smoother is a multi-state Gibbs sampler used to sample highly dimensional state variables that are correlated. The posterior draws from the above steps will be then used to compute posterior means and standard deviations.

We construct the time-varying coefficient of determination R_t^2 as

$$R_t^2 = 1 - \frac{\text{Var}(y_t - x_t'\beta_t)}{\text{Var}(y_t)},$$

where $\text{Var}(y_t - x_t'\beta_t) = x_t'P_tx_t + h^{-1}$, and $P_t = E[(\beta_t - \beta_{t-1})(\beta_t - \beta_{t-1})']$ is the covariance matrix of the state variables at time t conditional on information up to time $t - 1$ (prediction step of the Kalman filter). Note that without time-varying parameters, we would have $\text{Var}(y_t - x_t'\beta_t) = h^{-1}$, and hence the standard R^2 .

Table 1: The FX Market Share of UBS

The table displays the overall market share and the market share by customer type for UBS. The rank is with respect to the top 10 global leaders in the FX market from 2001 to 2011 based on the Euromoney annual survey. The market shares by customer type (available from 2003 for all customer types) are presented for *real money*, *leveraged funds* and *non-financial corporations*.

	Overall		Real		Leveraged		Non-financial	
	Market		Money		Funds		Corporations	
	share (%)	rank	share (%)	rank	share (%)	rank	share (%)	rank
2001	3.55	7	3.11	8	—	—	—	—
2002	10.96	2	10.77	2	—	—	—	—
2003	11.53	1	11.25	1	13.03	1	6.38	4
2004	12.36	1	11.32	2	11.70	2	7.16	3
2005	12.47	2	11.60	1	8.57	3	8.41	3
2006	22.50	1	11.35	2	5.23	7	6.38	4
2007	14.85	2	13.73	1	5.96	6	5.65	6
2008	15.80	2	9.07	2	7.53	4	5.13	5
2009	14.58	2	10.96	2	6.94	4	7.43	5
2010	11.30	2	9.39	2	14.63	2	4.93	9
2011	10.59	3	9.02	2	8.21	4	3.98	9

Table 2. Descriptive Statistics

The table presents descriptive statistics for daily log exchange rate returns and daily currency order flows. The exchange rate is defined as the US dollar price of a unit of foreign currency so that an increase in the exchange rate implies a depreciation of the US dollar. Order flow is defined as the difference between the value of buyer-initiated and seller-initiated transactions for the foreign currency so that a positive (negative) order flow implies net foreign currency purchases (sales). Order flows are in billions of US dollars and are classified into four customer segments: *asset managers*, *hedge funds*, *corporates* and *private clients*. Q_5 and Q_{95} are the 5th and 95th percentile, respectively. ρ_l is the autocorrelation coefficient for a lag of l trading days. The sample period ranges from January 2001 to May 2011. Exchange rates are from Thomson Reuters while customer order flows are proprietary data from UBS.

		<i>Mean</i>	<i>Sdev</i>	<i>Min</i>	<i>Max</i>	Q_5	Q_{95}	ρ_1	ρ_5	ρ_{21}
AUD	FX Returns (%)	0.0248	0.943	-7.627	8.219	-1.449	1.282	-0.077	-0.018	-0.040
	Asset Managers	-0.0017	0.146	-3.725	1.531	-0.151	0.144	-0.027	-0.006	-0.022
	Hedge Funds	-0.0052	0.118	-1.273	0.814	-0.175	0.154	0.073	0.041	-0.036
	Corporates	0.0035	0.048	-0.311	0.965	-0.035	0.050	0.172	0.109	0.041
	Private Clients	0.0002	0.092	-2.339	2.069	-0.067	0.072	-0.135	0.029	0.001
CAD	FX Returns (%)	0.0162	0.628	-3.298	3.770	-0.971	1.014	-0.027	-0.032	0.013
	Asset Managers	0.0035	0.136	-1.178	2.734	-0.136	0.149	0.096	0.010	0.022
	Hedge Funds	-0.0003	0.096	-0.754	1.162	-0.135	0.127	-0.007	0.020	-0.030
	Corporates	0.0048	0.056	-0.392	1.317	-0.042	0.051	0.172	0.106	0.034
	Private Clients	-0.0005	0.093	-4.023	1.043	-0.039	0.048	0.041	-0.021	-0.001
CHF	FX Returns (%)	0.0242	0.707	-2.873	5.038	-1.123	1.159	-0.058	-0.004	-0.052
	Asset Managers	-0.0033	0.199	-2.889	2.153	-0.266	0.228	0.032	0.048	0.005
	Hedge Funds	0.0091	0.207	-2.051	3.252	-0.261	0.288	-0.034	0.032	-0.009
	Corporates	0.0075	0.166	-5.702	3.572	-0.105	0.138	0.026	0.003	0.008
	Private Clients	0.0061	0.111	-1.348	2.597	-0.118	0.130	0.076	0.011	0.044
EUR	FX Returns (%)	0.0157	0.674	-3.173	3.733	-1.101	1.091	-0.022	0.009	-0.041
	Asset Managers	-0.0002	0.498	-12.803	3.981	-0.526	0.563	0.032	-0.001	-0.022
	Hedge Funds	-0.0267	0.391	-2.862	2.886	-0.590	0.580	-0.016	0.000	-0.008
	Corporates	-0.0490	0.166	-2.042	1.738	-0.296	0.169	-0.003	0.075	0.036
	Private Clients	0.0140	0.265	-2.122	4.240	-0.363	0.356	0.037	-0.004	0.000
GBP	FX Returns (%)	0.0036	0.618	-5.883	3.042	-0.986	0.949	0.026	-0.036	-0.037
	Asset Managers	0.0067	0.408	-8.289	9.102	-0.276	0.278	-0.130	0.024	0.019
	Hedge Funds	-0.0146	0.340	-13.162	3.183	-0.264	0.227	0.023	0.032	-0.004
	Corporates	0.0009	0.084	-0.914	1.815	-0.090	0.096	-0.009	0.046	-0.030
	Private Clients	0.0033	0.122	-1.698	1.321	-0.155	0.155	-0.004	0.006	-0.040
JPY	FX Returns (%)	0.0133	0.685	-6.203	3.706	-1.043	1.081	-0.054	0.013	-0.050
	Asset Managers	0.0090	0.306	-4.001	6.586	-0.326	0.329	0.127	-0.008	-0.008
	Hedge Funds	0.0127	0.280	-5.063	5.131	-0.327	0.352	-0.109	-0.001	-0.030
	Corporates	0.0050	0.061	-0.792	0.567	-0.078	0.089	0.037	-0.007	-0.030
	Private Clients	0.0004	0.102	-0.786	0.729	-0.144	0.136	0.014	0.014	-0.040
NOK	FX Returns (%)	0.0182	0.824	-4.709	5.625	-1.317	1.238	-0.019	-0.009	-0.051
	Asset Managers	0.0017	0.056	-0.638	0.605	-0.060	0.061	0.065	-0.037	0.014
	Hedge Funds	0.0001	0.040	-0.540	0.400	-0.051	0.050	0.074	0.051	0.024
	Corporates	0.0006	0.011	-0.112	0.127	-0.010	0.014	0.028	0.029	-0.006
	Private Clients	0.0003	0.010	-0.099	0.088	-0.012	0.012	0.061	-0.011	-0.004
NZD	FX Returns (%)	0.0233	0.911	-6.813	5.188	-1.546	1.351	-0.014	-0.015	-0.021
	Asset Managers	-0.0009	0.057	-1.171	0.672	-0.052	0.047	0.097	0.050	-0.012
	Hedge Funds	-0.0001	0.045	-0.440	0.633	-0.063	0.057	0.054	0.017	-0.007
	Corporates	-0.0014	0.015	-0.472	0.114	-0.014	0.010	0.183	-0.022	0.014
	Private Clients	-0.0001	0.019	-0.189	0.242	-0.023	0.026	0.062	0.004	-0.003
SEK	FX Returns (%)	0.0158	0.844	-5.379	5.243	-1.315	1.295	-0.029	0.007	-0.063
	Asset Managers	0.0001	0.057	-0.548	0.427	-0.078	0.081	-0.016	0.030	-0.016
	Hedge Funds	0.0006	0.044	-0.408	1.337	-0.044	0.045	0.032	0.068	-0.011
	Corporates	0.0005	0.018	-0.149	0.247	-0.018	0.020	0.049	-0.016	0.055
	Private Clients	0.0001	0.009	-0.102	0.145	-0.010	0.010	0.006	-0.038	0.021

Table 3. Cross-Correlations

The table shows the cross-correlations among daily log exchange rate returns and daily currency order flows. The exchange rate is defined as the US dollar price of a unit of foreign currency so that an increase in the exchange rate implies a depreciation of the US dollar. Order flow is defined as the difference between the value of buyer-initiated and seller-initiated transactions for the foreign currency so that a positive (negative) order flow implies net foreign currency purchases (sales). Order flows are in billions of US dollars and classified into four customer segments: *asset managers*, *hedge funds*, *corporates* and *private clients*. The superscripts *a*, *b*, and *c* denote statistical significance at the 10%, 5%, and 1% level, respectively. The sample period ranges from January 2001 to May 2011. Exchange rates are from Thomson Reuters while customer order flows are proprietary data from UBS.

		FX Returns	Asset Managers	Hedge Funds	Corporates	Private Clients
AUD	FX Returns	1.000				
	Asset Managers	0.061 ^c	1.000			
	Hedge Funds	0.200 ^c	-0.048 ^c	1.000		
	Corporates	-0.044 ^b	-0.008	-0.045 ^b	1.000	
	Private Clients	-0.051 ^c	-0.094 ^c	-0.087 ^c	0.022	1.000
CAD	FX Returns	1.000				
	Asset Managers	0.106 ^c	1.000			
	Hedge Funds	0.203 ^c	0.010	1.000		
	Corporates	-0.047 ^c	-0.071 ^c	-0.005	1.000	
	Private Clients	-0.092 ^c	-0.205 ^c	-0.225 ^c	-0.046 ^b	1.000
CHF	FX Returns	1.000				
	Asset Managers	0.149 ^c	1.000			
	Hedge Funds	0.312 ^c	0.003	1.000		
	Corporates	-0.072 ^c	-0.175 ^c	-0.041 ^b	1.000	
	Private Clients	-0.243 ^c	0.023	-0.110 ^c	0.047 ^b	1.000
EUR	FX Returns	1.000				
	Asset Managers	0.049 ^c	1.000			
	Hedge Funds	0.130 ^c	-0.186 ^b	1.000		
	Corporates	-0.056 ^c	0.038 ^b	-0.016 ^c	1.000	
	Private Clients	-0.348 ^c	-0.146	-0.027 ^c	0.010 ^c	1.000
GBP	FX Returns	1.000				
	Asset Managers	0.075 ^c	1.000			
	Hedge Funds	0.336 ^c	-0.039 ^c	1.000		
	Corporates	-0.079 ^c	-0.043 ^b	-0.089	1.000	
	Private Clients	-0.344 ^c	-0.007 ^c	-0.170	0.121	1.000
JPY	FX Returns	1.000				
	Asset Managers	0.103 ^c	1.000			
	Hedge Funds	0.227 ^c	0.022	1.000		
	Corporates	-0.050 ^c	-0.020	-0.009	1.000	
	Private Clients	-0.283 ^c	-0.115 ^c	-0.181 ^c	0.103 ^c	1.000
NOK	FX Returns	1.000				
	Asset Managers	0.068 ^c	1.000			
	Hedge Funds	0.083 ^c	0.011	1.000		
	Corporates	-0.030	-0.073 ^c	-0.074 ^c	1.000	
	Private Clients	0.147 ^c	0.016	0.048 ^b	-0.118 ^c	1.000
NZD	FX Returns	1.000				
	Asset Managers	0.114 ^c	1.000			
	Hedge Funds	0.132	-0.077 ^c	1.000		
	Corporates	0.013	-0.017	0.070 ^c	1.000	
	Private Clients	-0.014	-0.072 ^c	-0.023	0.036 ^a	1.000
SEK	FX Returns	1.000				
	Asset Managers	0.103 ^c	1.000			
	Hedge Funds	0.065 ^c	-0.079 ^c	1.000		
	Corporates	-0.007	-0.049 ^c	-0.027	1.000	
	Private Clients	0.086 ^c	0.032 ^a	0.066 ^c	-0.078 ^c	1.000

Table 4. The Predictive Ability of Order Flow

The table reports the in-sample and out-of-sample economic value of the predictive ability of monthly order flow. The results are based on dynamic asset allocation strategies investing in the G10 currencies with monthly rebalancing. The benchmark strategy is the naïve random walk (RW) model. The strategies condition on lagged currency order flow, which is classified into four customer segments: *asset managers* (AM), *hedge funds* (HF), *corporates* (CO) and *private clients* (PC). TOT indicates a strategy that conditions on total (aggregate) customer order flows. ALL is a strategy that conditions on all four (disaggregated) customer order flows. Using the exchange rate forecasts from each model, a US investor builds a maximum expected return strategy subject to a target volatility $\sigma_p^* = 10\%$ and proportional endogenous transaction costs. The strategy invests in a domestic bond and nine foreign bonds and is rebalanced monthly. For each strategy, we report the annualized mean (r_p), annualized volatility (σ_p), skewness ($Skew$), excess kurtosis ($Kurt$), annualized Sharpe ratio ($\mathcal{SR}$), annualized Sortino ratio ($\mathcal{SO}$), maximum drawdown ($\mathcal{MDD}$), and annualized performance fee ($\mathcal{P}$) a risk-averse investor is willing to pay to switch from the benchmark strategy to a competing strategy. $\mathcal{P}$ is the Goetzmann *et al.* (2007) premium return computed for $\gamma = 6$ and is expressed in annual basis points. The results are reported net of the effective bid-ask spread, which is assumed to be equal to 50% of the quoted spread. The in-sample period ranges from January 2001 to May 2011 and the out-of-sample period from January 2004 to May 2011. Exchange rates are from Thomson Reuters while customer order flows are proprietary data from UBS.

<i>Strategy</i>	r_p (%)	σ_p (%)	<i>Skew</i>	<i>Kurt</i>	ρ_1	$\mathcal{SR}$	$\mathcal{SO}$	$\mathcal{MDD}$ (%)	$\mathcal{P}$ (bps)
<i>In-Sample Period: Jan 2001 - May 2011</i>									
<i>RW</i>	9.3	10.9	-0.27	0.08	0.21	0.63	1.00	33.5	—
<i>AM</i>	16.4	11.5	-0.04	1.49	0.19	1.21	1.79	12.9	666
<i>HF</i>	17.0	11.2	0.23	0.83	0.02	1.30	2.40	10.4	753
<i>CO</i>	18.1	10.7	0.10	0.36	0.10	1.46	2.80	8.7	897
<i>PC</i>	11.0	11.0	-0.78	1.84	0.18	0.78	1.02	24.6	147
<i>TOT</i>	16.4	11.9	0.07	2.46	0.15	1.17	1.62	13.4	639
<i>ALL</i>	21.4	11.8	0.46	1.37	0.29	1.60	2.98	11.4	1158
<i>Out-of-Sample Period: Jan 2004 - May 2011</i>									
<i>RW</i>	5.3	13.4	-0.28	0.55	0.19	0.21	0.32	40.0	—
<i>AM</i>	13.4	12.9	0.06	0.56	0.16	0.84	1.46	21.1	879
<i>HF</i>	10.5	13.0	-0.28	0.62	0.14	0.61	0.92	29.9	560
<i>CO</i>	10.4	15.4	-0.23	0.47	0.24	0.51	0.81	33.0	345
<i>PC</i>	-5.0	17.2	-1.53	4.52	0.38	-0.44	-0.46	65.0	-1591
<i>TOT</i>	8.0	13.8	-1.18	3.11	0.06	0.40	0.46	28.4	190
<i>ALL</i>	13.2	13.5	-0.06	1.01	0.11	0.79	1.23	23.8	801

Table 5. Order Flow and Macroeconomic Information

The table presents regression results on the relation between monthly portfolio excess returns generated by conditioning out-of-sample on order flow and monthly portfolio excess returns by conditioning out-of-sample on macroeconomic information. The order flow strategies invest in the G10 currencies by conditioning on the order flow of asset managers (AM), hedge funds (HF), corporates (CO) and private clients (PC). The macroeconomic information strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011. The portfolio returns are computed net of the effective bid-ask spread, which is assumed to be equal to 50% of the quoted spread. The sign of the slopes of the predictive regressions used to generate the forecasts are constrained to be consistent with economic theory: all slopes are set to be positive except for FP and NXA. Asymptotic standard errors are in parentheses. The superscripts *a*, *b*, and *c* denote statistical significance at the 10%, 5%, and 1% level determined by bootstrapped *p*-values using 10,000 resamples of the portfolio weights by means of moving block bootstrap (see Gonçalves and White, 2005).

	α	β_{RW}	β_{FP}	β_{UIP}	β_{TR}	β_{PPP}	β_{MF}	β_{NXA}	β_{MOM}	$\overline{R}^2$
<i>Out-of-Sample Period: Jan 2004 - May 2011</i>										
AM	0.078 (0.052)	0.295 (0.141)				0.599 (0.280)	-0.116 (0.113)	-0.123 (0.173)	-0.100 (0.103)	16.5
	0.093 (0.054)			-0.231 (0.157)		0.834 ^b (0.245)	-0.075 (0.122)	-0.056 (0.167)	-0.134 (0.122)	12.2
	0.086 (0.057)				0.168 (1.943)	0.554 (1.692)	-0.106 (0.118)	-0.024 (0.164)	-0.159 (0.119)	8.6
	0.094 (0.053)	0.124 (0.182)	0.224 (0.174)	-0.228 (0.138)	0.099 (1.886)	0.522 (1.606)	-0.078 (0.088)	-0.191 (0.162)	-0.106 (0.113)	17.9
HF	0.045 (0.051)	0.276 (0.139)				0.633 ^b (0.241)	0.038 (0.120)	-0.150 (0.101)	-0.020 (0.099)	15.4
	0.047 (0.053)			0.278 (0.141)		0.579 ^b (0.237)	-0.002 (0.124)	-0.011 (0.084)	-0.098 (0.104)	13.9
	0.047 (0.060)				0.745 (1.031)	0.059 (0.872)	0.068 (0.124)	-0.072 (0.095)	-0.088 (0.127)	9.8
	0.022 (0.049)	0.323 (0.153)	0.014 (0.192)	0.376 ^a (0.160)	1.237 (1.317)	-0.724 (1.259)	0.021 (0.132)	-0.143 (0.127)	-0.073 (0.104)	22.3
CO	0.032 (0.049)	0.352 ^b (0.119)				1.240 ^c (0.188)	-0.095 (0.104)	-0.113 (0.169)	-0.240 (0.139)	40.2
	0.037 (0.053)			0.223 (0.173)		1.245 ^c (0.224)	-0.125 (0.131)	0.044 (0.186)	-0.327 (0.162)	34.7
	0.069 (0.051)				-2.513 ^a (0.947)	3.643 ^b (0.857)	-0.178 (0.131)	0.069 (0.166)	-0.251 (0.141)	39.9
	0.053 (0.045)	0.311 (0.150)	0.110 (0.170)	0.205 (0.137)	-2.144 (1.02)	2.995 ^a (0.969)	-0.204 (0.114)	-0.061 (0.153)	-0.216 (0.124)	48.4
PC	-0.102 (0.066)	0.738 ^a (0.204)				0.203 (0.314)	-0.155 (0.188)	0.120 (0.198)	-0.158 (0.129)	39.9
	-0.081 (0.079)			0.137 (0.226)		0.395 (0.374)	-0.166 (0.285)	0.398 (0.304)	-0.310 (0.197)	13.1
	-0.056 (0.065)				-2.108 (2.286)	2.379 (1.936)	-0.219 (0.316)	0.427 (0.309)	-0.251 (0.174)	16.6
	-0.085 (0.052)	0.697 ^a (0.234)	0.111 (0.199)	0.207 (0.183)	-1.707 (1.970)	1.560 (1.646)	-0.249 (0.242)	0.162 (0.211)	-0.144 (0.115)	44.1

**Table 6. The Predictive Ability of Combined Forecasts
Conditioning on Macroeconomic Information**

The table reports the out-of-sample economic value of the predictive ability of combined forecasts using macroeconomic information. The results are based on dynamic asset allocation strategies investing in the G10 currencies with monthly rebalancing. The benchmark strategy is the naïve random walk (RW) model. The competing strategies combine the *random walk* (RW), *forward premium* (FP), *uncovered interest parity* (UIP), *Taylor rule* (TR), *purchasing power parity* (PPP), *monetary fundamentals* (MF), *cyclical external imbalances* (NXA) and *momentum* (MOM). The combination employs the average (AVE), median (MED), trimmed mean (TRI), and mean-squared error (MSE) of the forecasts, the “kitchen sink” (KS) regression that incorporates all predictors into a multiple predictive regression, and Principal Component Analysis (PCA) that uses the first principal component. Using the exchange rate forecasts from each model combination, a US investor builds a maximum expected return strategy subject to a target volatility $\sigma_p^* = 10\%$ and proportional endogenous transaction costs. The strategy invests in a domestic bond and nine foreign bonds and is rebalanced monthly. For each strategy, we report the annualized mean (r_p), annualized volatility (σ_p), skewness ($Skew$), excess kurtosis ($Kurt$), annualized Sharpe ratio ($\mathcal{SR}$), annualized Sortino ratio ($\mathcal{SO}$), maximum drawdown ($\mathcal{MDD}$), and annualized performance fee ($\mathcal{P}$) a risk-averse investor is willing to pay to switch from the benchmark strategy to a competing strategy. $\mathcal{P}$ is the Goetzmann *et al.* (2007) premium return computed for $\gamma = 6$ and is expressed in annual basis points. The results are reported net of the effective bid-ask spread, which is assumed to be equal to 50% of the quoted spread. The out-of-sample period from January 2004 to May 2011. Exchange rates are from Thomson Reuters while customer order flows are proprietary data from UBS.

	r_p	σ_p	$Skew$	$Kurt$	$AR(1)$	$\mathcal{SR}$	$\mathcal{SO}$	$\mathcal{MDD}$	$\mathcal{P}$
<i>Strategy</i>	(%)	(%)						(%)	(bps)
<i>Out-of-Sample Period: Jan 2004 - May 2011</i>									
<i>RW</i>	5.3	13.5	−0.28	0.55	0.19	0.21	0.32	−40.05	
<i>AVE</i>	−2.7	19.6	−0.44	5.12	−0.30	−0.27	−0.36	−36.26	−1585
<i>MED</i>	1.4	20.3	−0.32	4.17	−0.18	−0.06	−0.08	−42.78	−1240
<i>TRI</i>	−2.8	18.2	−0.39	2.46	−0.29	−0.29	−0.42	−40.56	−1336
<i>MSE</i>	−2.1	22.1	−0.88	5.20	−0.16	−0.21	−0.26	−54.42	−2068
<i>KS</i>	−0.8	20.2	−0.22	3.45	−0.24	−0.17	−0.23	−34.20	−1412
<i>PCA</i>	−0.9	19.4	−0.65	1.58	−0.23	−0.18	−0.22	−46.66	−1316




Figure 1. Correlation between Currency Order Flow and FX Excess Returns across Horizons

This figure plots the contemporaneous correlation coefficient (vertical axis) between customer order flow and exchange rate returns against horizon (horizontal axis, log scale in days) for the G-10 currencies. The solid blue line is the average correlation coefficient across all exchange rates computed using overlapping return windows from 1 (10^0) to 252 ($> 10^2$) trading days. The dashed red lines represent the 90% bootstrap confidence intervals, estimated by generating 10,000 replications. Order flow is classified into four customer segments: *asset managers*, *hedge funds*, *corporates* and *private clients*. The sample period ranges from January 2001 to May 2011. Exchange rates are from Thomson Reuters while customer order flows are proprietary data from UBS.




Figure 2. Cumulative Wealth of Currency Order Flow Strategies

This figure displays the out-of-sample cumulative wealth of dynamic investment strategies conditioning on monthly currency order flow (solid red line) relative to the naïve random walk benchmark (dashed blue line). Order flow is classified into four customer segments: *asset managers*, *hedge funds*, *corporate clients* and *private clients*. The initial wealth is set at \$1, thereafter growing at the portfolio return, net of transaction costs. The out-of-sample analysis runs from January 2004 to May 2011. Exchange rates are from Thomson Reuters while customer order flows are proprietary data from UBS.




Figure 3. Order Flow and Macroeconomic Information: Time-Varying Alphas

The figure displays the time-varying alphas (α) on the relation between monthly currency order flow and macroeconomic information. The alphas are plotted in the blue line, while the red lines show the 95% confidence interval. The alphas are based on time-varying parameter regressions between monthly portfolio excess returns generated by conditioning out-of-sample on order flow and monthly portfolio excess returns by conditioning out-of-sample on macroeconomic information. The order flow strategies invest in the G-10 currencies by conditioning on the order flow of asset managers (AM), hedge funds (HF), corporates (CO) and private clients (PC). The macroeconomic information strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011. The portfolio returns are computed net of the effective bid-ask spread, which is assumed to be equal to 50% of the quoted spread.




Figure 4. Time-Varying Betas: Asset Managers

The figure displays the time-varying betas (β) and R^2 of the relation between monthly currency order flow and macroeconomic information for asset managers. The betas and R^2 are based on time-varying parameter regressions between monthly portfolio excess returns generated by conditioning out-of-sample on order flow and monthly portfolio excess returns generated by conditioning out-of-sample on macroeconomic information. The order flow strategies invest in the G-10 currencies by conditioning on the order flow of asset managers. The macroeconomic information strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011. The portfolio returns are computed net of the effective bid-ask spread, which is assumed to be equal to 50% of the quoted spread.




Figure 5. Time-Varying Betas: Hedge Funds

The figure displays the time-varying betas (β) and R^2 of the relation between monthly currency order flow and macroeconomic information for hedge funds. The betas and R^2 are based on time-varying parameter regressions between monthly portfolio excess returns generated by conditioning out-of-sample on order flow and monthly portfolio excess returns generated by conditioning out-of-sample on macroeconomic information. The order flow strategies invest in the G-10 currencies by conditioning on the order flow of hedge funds. The macroeconomic information strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011. The portfolio returns are computed net of the effective bid-ask spread, which is assumed to be equal to 50% of the quoted spread.




Figure 6. Time-Varying Betas: Corporates

The figure displays the time-varying betas (β) and R^2 of the relation between monthly currency order flow and macroeconomic information for corporates. The betas and R^2 are based on time-varying parameter regressions between monthly portfolio excess returns generated by conditioning out-of-sample on order flow and monthly portfolio excess returns generated by conditioning out-of-sample on macroeconomic information. The order flow strategies invest in the G-10 currencies by conditioning on the order flow of corporates. The macroeconomic information strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011. The portfolio returns are computed net of the effective bid-ask spread, which is assumed to be equal to 50% of the quoted spread.




Figure 7. Time-Varying Betas: Private Clients

The figure displays the time-varying betas (β) and R^2 of the relation between monthly currency order flow and macroeconomic information for private clients. The betas and R^2 are based on time-varying parameter regressions between monthly portfolio excess returns generated by conditioning out-of-sample on order flow and monthly portfolio excess returns generated by conditioning out-of-sample on macroeconomic information. The order flow strategies invest in the G-10 currencies by conditioning on the order flow of private clients. The macroeconomic information strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011. The portfolio returns are computed net of the effective bid-ask spread, which is assumed to be equal to 50% of the quoted spread.




Figure 8. Cumulative Order Flow for Macroeconomic Strategies: Asset Managers

The figure displays the cumulative monthly order flow of asset managers dedicated to a particular macroeconomic currency investment strategy (*blue line*) against the cumulative monthly wealth of the strategy (*red line*). The cumulative order flow is defined as the cumulative sum over time of the product of order flow for each currency times the weight for a currency implied by a macroeconomic strategy. It indicates the amount of money that asset managers would put into a macroeconomic strategy over time, and the figure relates it to the performance (cumulative wealth) of the strategy. The macroeconomic strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011.




Figure 9. Cumulative Order Flow for Macroeconomic Strategies: Hedge Funds

The figure displays the cumulative monthly order flow of hedge funds dedicated to a particular macroeconomic currency investment strategy (*blue line*) against the cumulative monthly wealth of the strategy (*red line*). The cumulative order flow is defined as the cumulative sum over time of the product of order flow for each currency times the weight for a currency implied by a macroeconomic strategy. It indicates the amount of money that hedge funds would put into a macroeconomic strategy over time, and the figure relates it to the performance (cumulative wealth) of the strategy. The macroeconomic strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011.




Figure 10. Cumulative Order Flow for Macroeconomic Strategies: Corporates

The figure displays the cumulative monthly order flow of corporates dedicated to a particular macroeconomic currency investment strategy (*blue line*) against the cumulative monthly wealth of the strategy (*red line*). The cumulative order flow is defined as the cumulative sum over time of the product of order flow for each currency times the weight for a currency implied by a macroeconomic strategy. It indicates the amount of money that corporates would put into a macroeconomic strategy over time, and the figure relates it to the performance (cumulative wealth) of the strategy. The macroeconomic strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011.




Figure 11. Cumulative Order Flow for Macroeconomic Strategies: Private Clients

The figure displays the cumulative monthly order flow of private clients dedicated to a particular macroeconomic currency investment strategy (*blue line*) against the cumulative monthly wealth of the strategy (*red line*). The cumulative order flow is defined as the cumulative sum over time of the product of order flow for each currency times the weight for a currency implied by a macroeconomic strategy. It indicates the amount of money that private clients would put into a macroeconomic strategy over time, and the figure relates it to the performance (cumulative wealth) of the strategy. The macroeconomic strategies are the following: the random walk (RW), forward premium (FP), uncovered interest parity (UIP), Taylor rule (TR), purchasing power parity (PPP), monetary fundamentals (MF), cyclical external imbalances (NXA) and momentum (MOM). All strategies are implemented out of sample for the period of January 2004 to May 2011.

References

- Akram, Q.F., D. Rime, and L. Sarno (2008). “Arbitrage in the Foreign Exchange Market: Turning on the Microscope,” *Journal of International Economics* **76**, 237–253.
- Albuquerque, R., E. de Francisco, and L. Marques (2008). “Marketwide Private Information in Stocks: Forecasting Currency Returns,” *Journal of Finance* **63**, 2297–2343.
- Bacchetta, P., and E. van Wincoop (2004). “A Scapegoat Model of Exchange-Rate Fluctuations?” *American Economic Review Papers and Proceedings* **94**, 114–118.
- Bacchetta, P., and E. van Wincoop (2006). “Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?” *American Economic Review* **96**, 552–576.
- Bacchetta, P., and E. van Wincoop (2011). “On the Unstable Relationship Between Exchange Rates and Macroeconomic Fundamentals,” Unpublished working paper, HEC Lausanne.
- Barroso, P., and P. Santa-Clara (2012). “Beyond the Carry Trade: Optimal Currency Portfolios,” Unpublished working paper, NOVA School of Business and Economics.
- Bank for International Settlements (2010). “Triennial Central Bank Survey: Report on Global Foreign Exchange Market Activity in 2010.”
- Bates, J.M., and C.W.J. Granger (1969). “The Combination of Forecasts,” *Operations Research Quarterly* **20**, 451–468.
- Bech, M. (2012). “FX Volume during the Financial Crisis and Now,” *BIS Quarterly Review* (March 2012), 33–43.
- Berger, D.W., A.P. Chaboud, S.V. Chernenko, E. Howorka, and J.H. Wright (2008). “Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data,” *Journal of International Economics* **75**, 93–109.
- Berger, D.W., A.P. Chaboud, E. Hjalmarsson, and E. Howorka (2009). “What Drives Volatility Persistence in the Foreign Exchange Market?” *Journal of Financial Economics* **94**, 192–213.
- Bjornes, G., and D. Rime (2005). “Dealer Behavior and Trading Systems in the Foreign Exchange Markets,” *Journal of Financial Economics* **75**, 571–605.
- Breedon, F., and P. Vitale (2010). “An Empirical Study of Portfolio-Balance and Information Effects of Order Flow on Exchange Rates,” *Journal of International Money and Finance* **29**, 504–524.
- Burnside, C., M. Eichenbaum, I. Kleshchelski, and S. Rebelo (2011). “Do Peso Problems Explain the Returns to the Carry Trade?” *Review of Financial Studies* **24**, 853–891.
- Burnside, C., M. Eichenbaum, and S. Rebelo (2009). “Understanding the Forward Premium Puzzle: A Microstructure Approach,” *American Economic Journal: Macroeconomics* **1**, 127–154.
- Campbell, J.Y., and S.B. Thompson (2008). “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?” *Review of Financial Studies* **21**, 1509–1531.
- Cheung, Y.-W., and M.D. Chinn (2001). “Currency Traders and Exchange Rate Dynamics: A Survey of the US Market,” *Journal of International Money and Finance* **20**, 439–471.
- de Jong, P., and N. Shephard (1995). “The Simulation Smoother for Time Series Models,” *Biometrika*, **82**, 339–350.

- Della Corte, P., L. Sarno, and G. Sestieri (2012). “The Predictive Information Content of External Imbalances for Exchange Rate Returns: How Much Is It Worth?,” *Review of Economics and Statistics* **94**, 100–115.
- Della Corte, P., L. Sarno, and I. Tsiakas (2009). “An Economic Evaluation of Empirical Exchange Rate Models,” *Review of Financial Studies* **22**, 3491–3530.
- Della Corte, P., L. Sarno, and I. Tsiakas (2011). “Spot and Forward Volatility in Foreign Exchange,” *Journal of Financial Economics* **100**, 496–513.
- Della Corte, P., L. Sarno, and I. Tsiakas (2012). “Volatility and Correlation Timing in Active Currency Management,” in James, J., L. Sarno and I.W. Marsh (eds.) *Handbook of Exchange Rates*. Hoboken, NJ: Wiley.
- Durbin, J., and S.J. Koopman (2012). *Time Series Analysis by State Space Methods*, Second Edition, Oxford University Press.
- Engel, C., N.C. Mark, and K.D. West (2008). “Exchange Rate Models Are Not as Bad as You Think,” *NBER Macroeconomics Annual 2007*, 381–441.
- Evans, M.D.D. (2010). “Order Flows and the Exchange Rate Disconnect Puzzle,” *Journal of International Economics* **80**, 58–71.
- Evans, M.D.D., and R.K. Lyons (2002). “Order Flow and Exchange Rate Dynamics,” *Journal of Political Economy* **110**, 170–180.
- Evans, M.D.D., and R.K. Lyons (2005). “Meese-Rogoff Redux: Micro-based Exchange Rate Forecasting,” *American Economic Review Papers and Proceedings* **95**, 405–414.
- Evans, M.D.D., and R.K. Lyons (2008). “How Is Macro News Transmitted to Exchange Rates?” *Journal of Financial Economics* **88**, 26–50.
- Evans, M.D.D., and D. Rime (2012). “Micro Approaches to Foreign Exchange Determination,” in James, J., L. Sarno and I.W. Marsh (eds.) *Handbook of Exchange Rates*. Hoboken, NJ: Wiley.
- Fama, E.F. (1984). “Forward and Spot Exchange Rates,” *Journal of Monetary Economics* **14**, 319–338.
- Fleming, J., C. Kirby, and B. Ostdiek (2001). “The Economic Value of Volatility Timing,” *Journal of Finance* **56**, 329–352.
- Froot, K.A., and T. Ramadorai (2005). “Currency Returns, Intrinsic Value, and Institutional-Investor Flows,” *Journal of Finance* **60**, 1535–1566.
- Goetzmann, W., J. Ingersoll, M. Spiegel, and I. Welch (2007). “Portfolio Performance Manipulation and Manipulation-Proof Performance Measures,” *Review of Financial Studies* **20**, 1503–1546.
- Gonçalves, S., and H. White (2005). “Bootstrap Standard Error Estimates for Linear Regression,” *Journal of the American Statistical Association* **100**, 970–979.
- Gourinchas, P.-O., and H. Rey (2007). “International Financial Adjustment,” *Journal of Political Economy* **115**, 665–703.
- Groen, J.J.J. (2000). “The Monetary Exchange Rate Model as a Long-Run Phenomenon,” *Journal of International Economics* **52**, 299–319.

- Hodrick, R.J., and E.C. Prescott (1997). “Postwar U.S. Business Cycles: An Empirical Investigation,” *Journal of Money, Credit and Banking* **29**, 1–16.
- King, M.R., C.L. Osler, and D. Rime (2012). “Foreign Exchange Market Structure, Players and Evolution,” in James, J., L. Sarno and I.W. Marsh (eds.) *Handbook of Exchange Rates*. Hoboken, NJ: Wiley.
- Koop, G. (2003). *Bayesian Econometrics*, New York: Wiley.
- Kyle, A.S. (1985). “Continuous Auctions and Insider Trading,” *Econometrica* **53**, 1315–1335.
- Lobo, M.S., M. Fazel, and S. Boyd (2007). “Portfolio Optimization with Linear and Fixed Transaction Costs,” *Annals of Operations Research* **152**, 341–365.
- Love, R., and R. Payne (2008). “Macroeconomic News, Order Flows, and Exchange Rates,” *Journal of Financial and Quantitative Analysis* **43**, 467–488.
- Lustig, H., N. Roussanov, and A. Verdelhan (2011). “Common Risk Factors in Currency Markets,” *Review of Financial Studies* **24**, 3731–3777.
- Lyons, R.K. (2001). *The Microstructure Approach to Exchange Rates*, Cambridge: MIT Press.
- Mancini-Griffoli, T., and A. Ranaldo (2011). “Limits to Arbitrage During the Crisis: Funding Liquidity Constraints and Covered Interest Parity,” Swiss National Bank working paper 2010-14.
- Mark, N.C. (2009). “Changing Monetary Policy Rules, Learning, and Real Exchange Rate Dynamics,” *Journal of Money, Credit and Banking* **41**, 1047–1070.
- Meese, R.A., and K. Rogoff (1983). “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?” *Journal of International Economics* **14**, 3–24.
- Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012a). “Carry Trades and Global Foreign Exchange Volatility,” *Journal of Finance* **67**, 681–718.
- Menkhoff, L., L. Sarno, M. Schmeling, and A. Schrimpf (2012b). “Currency Momentum Strategies,” *Journal of Financial Economics* **106**, 660–684.
- Molodtsova, T., and D.H. Papell (2009). “Out-of-Sample Exchange Rate Predictability with Taylor Rule Fundamentals,” *Journal of International Economics* **77**, 167–180.
- Osler, C.L., A. Mende, and L. Menkhoff (2011). “Price Discovery in Currency Markets,” *Journal of International Money and Finance* **30**, 1696–1718.
- Pojarliev, M., and R.M. Levich (2008). “Do Professional Currency Managers Beat the Benchmark?” *Financial Analysts Journal* **64**, 18–32.
- Pojarliev, M., and R.M. Levich (2010). “Trades of the Living Dead: Style Differences, Style Persistence and Performance of Currency Fund Managers,” *Journal of International Money and Finance* **29**, 1752–1775.
- Pojarliev, M., and R.M. Levich (2012). “Hunting for Alpha Hunters in the Currency Jungle,” *Journal of Portfolio Management* **39**, 3–7.
- Rapach, D.E., J.K. Strauss, and G. Zhou (2010). “Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy,” *Review of Financial Studies* **23**, 821–862.

- Rapach, D.E., and M.E. Wohar (2002). "Testing the Monetary Model of Exchange Rate Determination: New Evidence from a Century of Data," *Journal of International Economics* **58**, 359–385.
- Rime, D., L. Sarno, and E. Sojli (2010). "Exchange Rate Forecasting, Order Flow and Macroeconomic Information," *Journal of International Economics* **80**, 72–88.
- Sager, M.J., and M.P. Taylor (2006). "Under the Microscope: The Structure of the Foreign Exchange Market," *International Journal of Finance and Economics* **11**, 81–95.
- Stock, J.H., and M.W. Watson (2004). "Combination Forecasts of Output Growth in a Seven-Country Data Set," *Journal of Forecasting* **23**, 405–430.
- Taylor, J.B. (1993). "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* **39**, 195–214.
- Taylor, A.M., and M.P. Taylor (2004). "The Purchasing Power Parity Debate," *Journal of Economic Perspectives* **18**, 135–158.
- Timmermann, A. (2006). "Forecast Combinations," in Elliott, G., C.W.J. Granger, and A. Timmermann (eds.), *Handbook of Economic Forecasting*, Amsterdam: Elsevier.
- Welch, I., and A. Goyal (2008). "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction," *Review of Financial Studies* **21**, 1455–1508.