Currency Order Flow and Macroeconomic Information*

Pasquale DELLA CORTE
Imperial College London
p.dellacorte@imperial.ac.uk

Dagfinn RIME
Norges Bank & NTNU
dagfinn.rime@norges-bank.no

Lucio SARNO
Cass Business School & CEPR
lucio.sarno@city.ac.uk

Ilias TSIAKAS
University of Guelph
itsiakas@uoguelph.ca

July 2012

*__Acknowledgements:__ The authors are indebted for constructive comments to Rui Albuquerque, Philippe Bacchetta, Ekkehart Boehmer, Nicola Borri, Giuseppe De Arcangelis, Charles Jones, Nengjiu Ju, Michael King, Robert Kosowski, Michael Moore, Marco Pagano, Alessandro Palandri, Lasse Pedersen, Tarun Ramadorai, Thomas Stolper, Adrien Verdhelan, Paolo Vitale, Kathy Yuan, 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, and the 2012 SIRE Econometrics Workshop in Glasgow. 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 Science and Humanities Research Council of Canada. **Corresponding author:** Lucio Sarno, Cass Business School, City University, 106 Bunhill Row, London EC1Y 8TZ, UK.

Currency Order Flow and Macroeconomic Information

Abstract

This paper investigates empirically whether currency order flow aggregates disperse public information about economic fundamentals that are relevant to exchange rates. Our analysis uses a unique data set on end-user transactions across four customer groups for the G10 currencies from 2001 to 2011. We find that customer order flow has substantial predictive ability for exchange rate returns leading to highly profitable trading strategies net of transaction costs. More importantly, a large part of the predictive information of order flow can be explained by a time-varying combination of macroeconomic fundamentals. Furthermore, there is no evidence of a significant "alpha" in customer order flow as it contains no additional predictive information over and above widely available macroeconomic information.

Keywords: Exchange Rates; Order Flow; Market Microstructure; Forecasting; Asset Allocation.

JEL Classification: F31; G11; G15.

1 Introduction

Does currency order flow aggregate disperse macroeconomic information across different market participants? This question is at the center of the market microstructure approach to exchange rates pioneered by Evans and Lyons (2002). The microstructure 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 what is widely known as the "exchange rate disconnect" puzzle (Obstfeld and Rogoff, 2001). The market microstructure literature asserts that transactions can affect prices because they convey information. News can be impounded directly in currency prices or indirectly via order flow (Evans and Lyons, 2008).[1] Order flow can also affect the price for reasons unrelated to publicly available news (e.g., changing risk aversion, liquidity and hedging demands).

This paper investigates empirically the predictive ability of customer order flow and its relation to macroeconomic information using a new and comprehensive order flow data set obtained from UBS, a global leader in Foreign Exchange (FX) trading. The data set disaggregates customer order flows into trades executed between UBS (the dealer) and four segments (customers): asset managers, hedge funds, corporates and private clients. Overall, this is a rich data set that contains the US dollar value of disaggregated daily order flows over a sample period ranging from January 2001 to May 2011 and covers the G10 currencies. Furthermore, all macroeconomic variables are constructed using real-time data that was available to market participants at the time forecasts are made. These data, therefore, provide us with a unique opportunity to examine the predictive ability of customer order flow and it relation to real-time macroeconomic fundamentals over a long sample and a large set of exchange rates.

Armed with these data, our paper addresses three main questions. First, can customer order flow predict exchange rate returns? We answer this question from the point of view of an investor (or dealer) implementing a dynamic asset allocation strategy across the G10 currencies. We choose a trading strategy to assess the predictive ability of customer order flow in order to measure the tangible economic gains of predictability and because it is through trades that customers reveal their information.

[1] Evans and Lyons (2008) provide an excellent example of the indirect channel to the price adjustment process. Consider a scheduled macro announcement on US GDP growth that is higher than the expectation of market participants. Suppose that everyone agrees that the GDP announcement represents good news for the US dollar but there are diverse opinions as to how large the appreciation should be relative to the Japanese yen. In this case, some participants may view the initial rise in the yen/dollar spot rate as too large while others as too small. Those who view the rise as too small will place orders to purchase the dollar, while those who view the rise as too large will place orders to sell. Positive order flow signals that the initial yen/dollar spot rate was below the balance of opinion among market participants and vice versa.

Second, can macroeconomic information explain the predictive ability of order flow? For this question, we relate the portfolio returns generated ex post by conditioning on customer order flow to the portfolio returns generated ex post 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 changes in interest rates, real exchange rates or monetary fundamentals and the extent to which they reflect information not related to economic fundamentals. We then ask a complementary question: can forecast combinations conditioning on macroeconomic variables replicate ex ante the predictive ability of order flow? 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, however, it could be that order flow summarizes the available macroeconomic information in a distinct and effective manner that cannot be replicated by a standard forecast combination.

Finally, third, does the relation between order flow and macroeconomic information vary 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, 2006), where every day the market may focus its attention on a different macroeconomic variable (the scapegoat). The scapegoat theory relies on traders assigning a different weight to a macroeconomic indicator every day as the market rationally searches for an explanation for the observed exchange rate change.

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. The trades between dealers and customers are not transparent, since 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 and future risk premia. This information is then revealed to the rest of the market when dealers trade with each other motivated primarily by liquidity and inventory concerns.[2] Trades between dealers and customers account for 61.1% of FX turnover (Bank for International Settlements, 2010).[3]

[2] In 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 Systems. 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 only available to dealers (see, e.g., Bjonnes and Rime, 2005).

[3] For further details on the institutional structure of the FX market, see, for example, Lyons (2001), Bjonnes and Rime (2005), Evans and Lyons (2006), Sager and Taylor (2006), and Evans (2011).

This trading mechanism implies that customer order flow may be a predictor of future FX excess returns. Order 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.[4] The argument is as follows (see, e.g., Evans and Lyons, 2005, 2006, 2007). The spot exchange rate is the rate quoted by FX dealers and hence reflects the dealers' information set. If dealers first receive the information conveyed by customer order flow and subsequently incorporate it in their quotes, then customer order flow should be able to forecast FX excess returns. Note that the information conveyed by customer order flow can only be used by the dealer who facilitated the transaction as it is not observed by other market participants. Furthermore, customers are heterogeneous in their motivation for trading, attitude towards risk and horizon leading them to adopt different trading strategies. Therefore, different customer groups will provide dealers with different information. Through interdealer trading this information will be aggregated and mapped to a price thus establishing a transmission mechanism from customer order flow to the exchange rate.

We find that customer order flow for currencies has substantial predictive 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 predictability, especially at the monthly horizon. More importantly, combinations of empirical exchange rate models based on macroeconomic information can explain *ex post* a large part (up to 50%) of the predictive ability of order flow. These exchange rate models include the random walk, forward premium, uncovered interest parity, purchasing power parity, monetary fundamentals, Taylor rule, cyclical external imbalances and momentum. Furthermore, there is no evidence of a significant "alpha" in customer order flow. In other words, there is no additional predictive information in order flow over and above the information embedded in macroeconomic fundamentals. Standard forecast combinations of macroeconomic variables fail to replicate *ex ante* the predictive ability of order flow. This leads us to conclude that order flow provides a distinct and effective way of aggregating macroeconomic information.

Finally, 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 different macroeconomic fundamentals. This is particularly evident before versus after the crisis as, for example, after the crisis the carry trade is 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, 2006). By allowing for time-variation in the relation between order flow and macroeconomic information, the latter can on average explain 50% to 70% of the former. To conclude, overall we interpret this evidence as

[4] 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).

suggesting that the predictive information content in order flow is not only economically important but also derives from aggregating disperse public information about economic fundamentals that are relevant to exchange rates.

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. 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 Section 6 concludes.

2 Data and Preliminaries

2.1 Exchange Rate Data

The empirical analysis uses spot and forward exchange rates, customer order flows, interest rates 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 kroner (NOK), New Zealand dollar (NZD) and Swedish kronor (SEK). The data range from January 2001 to May 2011 and cover 2618 daily observations after removing holidays and weekends or equivalently 125 monthly observations.

The exchange rates are Thomson Reuters data obtained through *Datastream*. For the daily analysis, we use daily spot and spot-next forward rates, whereas for the monthly analysis we use end-of-month spot and one-month forward rates. 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. Most of the empirical work uses mid-quotes, but bid and ask quotes are used to construct transaction costs.

2.2 Order Flow Data

The order flow data come from proprietary daily transactions between four end-user segments (customer groups) and UBS, a global leader in the FX market. Order flows are disaggregated into four segments: trades executed between UBS and asset managers (AM), hedge funds (HF), corporates (CO) and private clients (PC). The asset managers segment comprises long-term 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. The corporates segment includes non-financial corporations that import or export products and services around the world or have an international supply chain. Treasury units of large non-financial corporations are treated as corporates unless they pursue an aggressive (highly leveraged) investment strategy, in which case they are classified as hedge funds. The final segment, private clients, includes wealthy clients with in excess of \$3 million in investible liquid assets. Private clients trade primarily for financial reasons and with their own money.

Table 1 reports UBS' market share by customer type and its rank relative to the top 10 global FX dealers from 2001 to 2011 based on the Euromoney annual survey.[5] The table reveals that UBS has been one of the top dealers for both the overall market and particular end-user segments. Although the Euromoney survey uses different groupings than UBS, three of the groups defined by Euromoney (real money, leveraged funds and non-financial corporations) seem to align well with three of the UBS segments (asset managers, hedge funds and corporates).[6] 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 with non-financial corporations.[7] The latest estimate of the average daily turnover in FX is \$4.7 trillion (Bech, 2012). As UBS has approximately 10% overall market share, then it must produce a daily turnover of more than \$400 billion per day. This figure rises to more than \$8 trillion per month from UBS alone.

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 UBS clients. 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 (negative) sign if the initiator of the transaction (the non-quoting counterparty) is buying (selling) foreign currency. 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 the most comprehensive in this literature to date and is also 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. Second, our data set spans more than 10 years of daily observations for nine

[5] See also Table C1 in the separate Appendix for a list of the top 10 global leaders in terms of market share in annual FX turnover from 2001 to 2011 based on the Euromoney FX Survey.

[6] Euromoney also have 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, but these "customer-banks" often have non-financial customers behind them.

[7] Based on our private discussions with traders, UBS is also one of the leading counterparty traders with private clients.

[8] It 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 by a customer acting on a dealer's quotes. Then transaction volume is 10, but order flow is –10 (see, e.g., Lyons, 2001).

currency pairs and it comes from a major FX market leader. Although there are recent studies that employ customer order flow data, they typically suffer from a number of 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, 2006, 2007) and Evans (2010) employ six years of data for one currency pair from Citibank. Cerrato, Sarantis and Saunders (2011) use six years of data for nine currency pairs from UBS but, in contrast to this paper, their data are weekly; they have 317 weekly observations compared to our 2618 daily observations. 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 with 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). Finally, our order flow data are raw data with minimal filtering. For instance, data are adjusted to take into account large merger and acquisition deals which are announced weeks or months in advance. Cross-border merger 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-publisized and thus are anticipated in advance by market participants. Furthermore, 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 confidentiality issues, flows from prop traders because they trade with UBS' own money, while small banks often have non-financial customers behind them.

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. 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 collect and compile real-time economic data from historical paper and electronic sources for the following countries: Australia, Canada, 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] A list of the data sources is reported below.

[9] For example, as of 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

For industrial production, we obtain real-time seasonally adjusted data on industrial production indices through the OECD *Main Economic Indicators*. These are available at monthly frequency, except for Australia, New Zealand and Switzerland which are quarterly. For real output, we collect real-time seasonally adjusted data on gross domestic product from the *OECD Main Economic Indicators*. These data are available in national currency at quarterly frequency for all countries. For the output gap, we construct the deviations from the Hodrick and Prescott (1997) filter as in Molodtsova and Papell (2009). Note that we update the Hodrick and Prescott (1997) trend each period so that ex-post data is not used to construct the output gap. In other words, at time t we only use data up to $t-1$ to construct the output gap.[10] For consumer prices, we collect real-time seasonally adjusted observations on consumer price indices from the *OECD Main Economic Indicators*. Data are published every month for all countries, except Australia and New Zealand for which data are available quarterly. For money supply, we collect real-time seasonally adjusted data on broad money from the OECD *Main Economic Indicators*. 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*. 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*. Data are published every quarter in national currency for all countries. Finally, for interest rates, we use daily (end-of-month) spot-next (one-month) Eurodeposit rates from *Datastream*.

We convert the data by taking logs, except for interest rates and order flows. Henceforth the symbols s_t, f_t, x_t, i_t, m_t, π_t, y_t and $\overline{y}_t$ refer to the log spot exchange rate, log forward exchange rate, order flow, interest rate, log money supply, inflation rate, log real output and log output gap, respectively. We use an asterisk to denote the data (i_t^*, m_t^*, π_t^*, y_t^* and $\overline{y}_t^*$) for the foreign country.

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.

2004.

[10] The output gap for the first period is computed using real output data from January 1990 to January 2001. In the Hodrick and Prescott (1997) filter, we use a smoothing parameter equal to 14,400 as in Molodtsova and Papell (2009).

Table 3 presents the contemporaneous cross-correlations between daily customer order flows and returns. While asset managers and hedge funds are positively correlated with exchange rate returns, corporates and private clients are typically 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 interesting preliminary finding which indicates that different types of order flow may be predicting exchange rates to move in opposite directions. More generally, these results are consistent with previous empirical evidence reported by Evans and Lyons (2002) and Sager and Taylor (2006) indicating 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 types of macroeconomic fundamentals consistent with widely used empirical exchange rate models. All predictive regressions have the following linear structure:

$$\Delta s_{t+1} = \alpha + \beta x_t + \varepsilon_{t+1}, \tag{1}$$

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, α and β are constant parameters to be estimated, and ε_{t+1} is a normal error term. The predictive regressions differ only in the way they specify the predictive variable x_t that is used to forecast exchange rate returns.

Note that for the macroeconomic models we impose constraints on the sign of the slope parameter β in order to be consistent with the economic theory that each of these models represents. We will specify the constraints below. For example, Campbell and Thompson (2008) impose similar constraints on the sign of the slope of predictive regressions in assessing predictability in stock returns. For the order flow regressions, we impose no constraints.[11]

3.1 Models Conditioning on Order Flow

We estimate three types of predictive regressions that condition on customer order flow. The first type conditions separately on the order flow of each of the four customers, asset managers (x_t^{AM}), hedge

[11] The constraints fix the sign of the slope of the predictive regression to be the same as the sign expected by theory. Hence 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.

funds (x_t^{HF}), corporates (x_t^{CO}), and private clients (x_t^{PC}), leading to four regressions. The second type conditions on all four customer order flows in one regression, which we call the disaggregated order flows: $x_t = \left\{x_t^{AM}, x_t^{HF}, x_t^{CO}, x_t^{PC}\right\}$. Finally, we condition on the sum of the four order flows, which we call the aggregate (or total) order flow: $x_t = x_t^{AM} + x_t^{HF} + x_t^{CO} + x_t^{PC}$. These regressions will determine whether there is predictive information in customer order flow and the extent to which different customer groups convey different information.

3.2 Models Conditioning on Macroeconomic Information

3.2.1 Random Walk

The first specification based on public information is the driftless (or naive) random walk (RW) model that sets $\alpha = \beta = 0$. Since the seminal contribution of Meese and Rogoff (1983), this model has become the benchmark in assessing exchange rate predictability. The RW model captures the prevailing view in international finance research that exchange rates are unpredictable and forms the basis of the widely used carry trade strategy in active currency management (e.g., Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011; Lustig, Roussanov and Verdelhan, 2011; Menkhoff, Sarno, Schmeling and Schrimpf, 2012). The RW model is the benchmark to which we compare the predictive regressions conditioning on order flow.

3.2.2 Forward Premium and Uncovered Interest Parity

The second and third specifications use the forward premium (FP) as a 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 predictive regression using FP as conditioning information captures deviations from the uncovered interest rate parity (UIP) condition. Under risk neutrality and rational expectations, UIP implies that $\alpha = 0$, $\beta = 1$, and the error term is serially uncorrelated. However, empirical studies consistently reject the UIP condition and it is a stylized fact that estimates of β often display a negative sign (e.g., Evans, 2011, Ch. 11). This implies that high-interest rate currencies tend to appreciate rather than depreciate over time.[12] In this context, we henceforth call UIP the regression that imposes a positive sign on the slope of the forward premium and FP the regression that imposes a negative sign.

[12] Note that we 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. In this case, testing UIP is equivalent to testing for forward unbiasedness in exchange rates (Bilson, 1981). There is ample empirical evidence that CIP holds in practice for the data frequency examined in this paper. For recent evidence, see Akram, Rime and Sarno (2008). The only exception in our sample is the period following Lehman's bankruptcy, when the CIP violation persisted for a few months (e.g., Mancini-Griffoli and Ranaldo, 2011).

3.2.3 Purchasing Power Parity

The fourth regression is based on the purchasing power parity (PPP) condition and sets

$$x_t = p_t - p_t^* - s_t, \tag{3}$$

where p_t (p_t^*) is the log of the domestic (foreign) price level. This is equivalent to a trading strategy that buys undervalued currencies and sells overvalued currencies relative to PPP. The PPP hypothesis states that national price levels should be equal when expressed in a common currency and is typically thought of as a long-run condition rather than holding at each point in time (e.g., Rogoff, 1996; and Taylor and Taylor, 2004). In the PPP regression, we impose a positive sign on β.

3.2.4 Monetary Fundamentals

The fifth regression conditions on monetary fundamentals (MF):

$$x_t = (m_t - m_t^*) - (y_t - y_t^*) - s_t, \tag{4}$$

where m_t (m_t^*) is the log of the domestic (foreign) money supply and y_t (y_t^*) is the log of the domestic (foreign) real output. 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 in what is commonly referred to as the "exchange rate disconnect" puzzle (Obstfeld and Rogoff, 2001). On the other hand, there is growing evidence that exchange rates and monetary fundamentals are cointegrated, which requires that the exchange rate and/or the fundamentals move in a way to restore and equilibrium relation between them in the long run (e.g., Groen, 2000; Rapach and Wohar, 2002). In the MF regression, we impose a positive sign on β.

3.2.5 Taylor Rule

The sixth specification uses the Taylor (1993) rule (TR) defined as

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

where π_t (π_t^*) is the domestic (foreign) inflation rate, and $\overline{y}_t$ ($\overline{y}_t^*$) is the domestic (foreign) output gap measured as the percent deviation of real output from an estimate of its potential level computed using the Hodrick and Prescott (1997) filter.[13] The Taylor rule postulates that the central bank

[13] Note that 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.

raises the short-term nominal interest rate when output is above potential output and/or inflation rises above its desired level. 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, 2007; Mark, 2009; Molodtsova and Papell, 2009). In the TR regression, we impose a positive sign on β.[14]

3.2.6 Cyclical External Imbalances

The seventh model employs as the predictive variable a bilateral measure of cyclical external imbalances between the US and the foreign country. Following Gourinchas and Rey (2007), we construct nxa_t, a global measure of cyclical external imbalances, which linearly combines detrended (log) exports, imports, foreign assets, and liabilities relative to GDP. The bilateral measure of cyclical external imbalances between the US and a foreign country 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. In the NXA regression, we impose a negative sign on β, which is justified as follows: a country running external imbalances ($nxa_t < 0$) will experience a future currency depreciation ($\Delta s_{t+1} > 0$) 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).[15]

3.2.7 Momentum

The eighth and final specification uses the one-month rolling exchange rate return as the conditional mean of the one-period ahead exchange rate. This momentum (MOM) strategy produces a long exposure to the currencies that are trending higher, and a short exposure to the currencies that are trending lower. In the MOM regression, we impose a positive sign on β.

[14] We 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.

[15] Following Gourinchas and Rey (2007) closely, we filter out the trend component in (log) exports, imports, foreign assets, and liabilities relative to GDP using the Hodrick-Prescott filter. We then combine these stationary components with weights reflecting the (trend) share of exports and imports in the trade balance, and the (trend) share of foreign assets and liabilities in the net foreign assets, respectively. These time-varying weights are replaced with their sample averages to minimize the impact of measurement error. Finally, note that the Hodrick-Prescott filter and the constant weights are based on the full-sample information for the in-sample analysis. In the out-of-sample analysis, however, we implement the Hodrick-Prescott filter and compute the weights only using information available at the time of the forecast in order to avoid any look-ahead bias.

3.3 Combined Forecasts

In addition to specifying predictive regressions conditioning on individual macroeconomic fundamentals, we also combine the forecasts arising from the full set of empirical exchange rate models. Although the potentially superior performance of combined forecasts is known since the seminal work of Bates and Granger (1969), applications in finance are only recently becoming increasingly popular (Timmermann, 2006; Rapach, Strauss and Zhou, 2010).

Our empirical analysis estimates N predictive regressions for a vector of K 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 $\Delta s_{c,t+1|t}$ for the vector of exchange rate returns as the weighted average of the N individual forecasts:

$$\Delta s_{c,t+1|t} = \sum_{j=1}^{N} \theta_{j,t} \Delta s_{jt+1|t}, \tag{6}$$

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 (6); the median (MED) combination forecast is the median of $\left\{\Delta s_{j,t+1|t}\right\}_{j=1}^{N}$; and the trimmed (TRI) mean combination forecast sets $\theta_{j,t} = 0$ for the individual forecasts with the smallest and largest values and $\theta_{j,t} = 1/\left(N-2\right)$ for the remaining individual forecasts in Equation (6). 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 Watson (2004), and uses statistical information on the past performance of each individual model. In particular, it sets the weights by computing the following mean squared error (MSE) forecast combination:

$$\theta_{j,t} = \frac{MSE_{j,t}^{-1}}{\sum_{j=1}^{N} MSE_{j,t}^{-1}}, \quad MSE_{j,t} = \frac{1}{T}\sum_{t=1}^{T} \left(\Delta s_{j,t} - \Delta s_{j,t|t-1}\right)^2. \tag{7}$$

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

$$\Delta s_{t+1} = \alpha + \sum_{j=1}^{N} \beta_j x_t^j + \varepsilon_{t+1}. \tag{8}$$

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 macro-based forecast combinations to the RW and the order flow models. The objective of this exercise is to assess whether order flow conveys predictive information that can be captured by standard methods of combining macroeconomic information. An important advantage of these forecast combination methods is that they produce ex ante forecasts that can be used in realistic trading strategies.

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, Deutsche mark\euro, British pound, Japanese yen, Norwegian kroner, New Zealand dollar and Swedish kronor. 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, 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 $r_{t+1|t} = E_t\left[r_{t+1}\right]$ is the conditional expectation of r_{t+1} and $\Delta s_{t+1|t} = E_t\left[\Delta s_{t+1}\right]$ is the conditional expectation of Δs_{t+1}. Hence the only risk the US investor is exposed to is FX risk.

Every 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:

$$\begin{aligned} \max_{w_t} \quad r_{p,t+1|t} &= w_t' r_{t+1|t} + (1 - w_t'\iota) r_f - \phi_{t+1|t} \\ \text{s.t.} \quad \sigma_p^* &= \left(w_t' V_{t+1|t} w_t\right)^{1/2}, \end{aligned} \tag{9}$$

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 $r_{p,t+1}$, 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:

$$\phi_{t+1|t} = \sum_{i=1}^{K} \tau_{i,t+1|t} \left| w_{i,t} - w_{i,t}^{-} \right|, \tag{10}$$

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:

$$er_{t+1}^{net} = s_{t+1}^{b} - f_t^{a}, \tag{11}$$

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 the 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:

$$\begin{aligned} er_{t+1}^{net} &= \left(s_{t+1} - \frac{s_{t+1}^{a} - s_{t+1}^{b}}{2}\right) - \left(f_t + \frac{f_t^{a} - f_t^{b}}{2}\right) \\ &= (s_{t+1} - f_t) - c_{t+1}, \end{aligned} \tag{12}$$

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 need to compute the time-varying weights w_t using information up to time t. These weights will 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.[16] Once the optimal weights are computed, the return on the investor's portfolio net of the realized transaction costs is equal to:

$$r_{p,t+1} = w_t' r_{t+1} + \left(1 - w_t'\iota\right) r_f - \phi_{t+1}. \tag{13}$$

Finally, note that we assume that $V_{t+1|t} = \overline{V}$, where $\overline{V}$ is the unconditional covariance matrix of exchange rate returns. In other words, we do not model the dynamics of FX return volatility and correlation. Therefore, the optimal weights will vary across the empirical exchange rate models only to the extent that the predictive regressions produce better forecasts of the exchange rate returns.[17]

4.3 Performance Measures

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

$$M\left(r_p\right) = \frac{1}{(1-\gamma)} \ln \left\{ \frac{1}{T} \sum_{t=1}^{T} \left(\frac{1 + r_{p,t}}{1 + r_f} \right)^{1-\gamma} \right\}, \tag{14}$$

where $M\left(r_p\right)$ 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 :

$$\mathcal{P} = M(r_p^*) - M(r_p^b), \tag{15}$$

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

[16] We 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).

[17] See Della Corte, Sarno and Tsiakas (2009) for an empirical analysis of the effect of dynamic volatility on mean-variance strategies in FX.

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*).[18]

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 only 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 fund.

5 Empirical Results

5.1 The Correlation between FX Returns and Currency Order Flow

We begin our empirical analysis by calculating the correlation between contemporaneous exchange rate returns and currency order flows 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 FX returns and order flows at the daily, monthly and longer horizons. For example, a significant positive correlation at the daily horizon suggests that daily order flow can explain daily FX returns. Furthermore, a significant positive correlation at the annual horizon makes it empirically plausible for order flow to have predictive power for future FX returns.

Figure 1 displays for each customer type the average correlation across all exchange rates. 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 $10,000$ replications under the null hypothesis that each order flow and FX return series is independent and identically distributed (i.i.d.). 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. For HF flows, the correlations are generally positive and higher than AM but less significant. For CO flows, the initial 1-day correlations are negative and insignificant, but become positive for long horizons. The PC flows are

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

negative for all horizons but highly statistically significant only for short horizons.

Overall, this preliminary analysis provides strong evidence that the order flow of different customer groups can have (positive or negative) contemporaneous correlation to FX returns which may extend to relatively long horizons.[19] Next we turn to exploring whether the information content in order flow has predictive power for future FX excess 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 dynamic asset allocation strategies that condition 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 through trades that customers reveal their information to dealers. Furthermore, this allows us to measure the tangible economic gains from exploiting the predictability of different types of customer order flow. In particular, we analyze the performance of dynamically rebalanced portfolios based on the order flow strategies relative to the random walk (RW) benchmark. In this setting, the investor obtains forecasts of exchange rate returns for next period (day or 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 $\gamma = 6$.[20]

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 set ranging from January 2001 to May 2011. For the out-of-sample analysis, we first estimate the predictive regressions over the initial sample period of January 2001 to December 2003, and 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 drawdown ($\mathcal{MDD}$) 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

[19] This preliminary evidence also effectively summarizes the contemporaneous regression results reported later in the separate Appendix.

[20] The 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.

taken into account in the optimization that delivers the mean-variance weights. We consider an effective transaction cost that is equal to 50% of the quoted spread.[21]

It is important to recall that customer transactions are private and the details (e.g., bid and ask quotes) are known to the dealer customers transact with but not to the rest of the market. 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 trading strategies 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.

We begin our discussion with the daily rebalancing results in Table 4. Our first finding is that the benchmark RW strategy has not performed particularly well in sample and especially out of sample. This is not surprising given that the crisis period beginning in June 2007 saw a collapse of carry trade strategies. In the context of a longer sample than the one used in this paper, the carry trade losses that characterize the 2007-2008 period would have a much smaller impact on average carry trade returns.[22] At any rate, the in-sample Sharpe ratio of the RW is 0.13, while the out-of-sample exercise leads to a negative Sharpe ratio of -0.17. It is also noteworthy that the $\mathcal{MDD}$ of the carry trade is very large: 37% in sample and 48% out of sample.

Turning to the out-of-sample evaluation of order flow models, the following findings are noteworthy. The order flow models outperform the RW benchmark in all cases thus indicating that order flow has substantial predictive information. The best out-of sample performing model is PC, which has a Sharpe ratio net of transaction costs of 0.40 (compared to -0.17 for RW) and a performance fee of 868 annual basis points (bps). The HF model also performs well with $\mathcal{SR} = 0.27$ and $\mathcal{P} = 759$ bps. AM is only type of order flow that produces a negative Sharpe ratio ($\mathcal{SR} = -0.03$), which however is still better than that of RW. In conclusion, these results show that a risk-averse investor will 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 ranges from 315 to 868 basis points for daily rebalancing. This is strong evidence of the predictive power of order flow information as compared to the RW benchmark.[23]

In interpreting these results, there seems to be a natural explanation of why the PC order

[21] It 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., Goyal and Saretto, 2009).

[22] For an evaluation of the carry trade performance using longer samples see, for example, Burnside, Eichenbaum, Kleshchelski and Rebelo (2011), Lustig, Roussanov and Verdelhan (2011), and Menkhoff, Sarno, Schmeling and Schrimpf (2012).

[23] In Tables C5 and C6 of the separate Appendix we report results for the same exercise where in the first case we account for serial correlation in order flow, and in the second case the predictive regression is estimated using the M-estimator, which is robust to outliers. The M-estimator is discussed in the separate Appendix. The results in Tables C5 and C6 confirm the superior performance of the order flow models and show that in some cases (e.g., AM) it improves significantly when using the M-estimator.

flow strategy performs best. Private clients are thought of as liquidity providers, who are rather uninformed and unsophisticated relative to asset managers and hedge funds. In fact, the sign of the PC order flow coefficient in the predictive regression is negative. This is consistent with the correlation illustrated in Figure 1. More importantly, it implies that the PC trading strategy exploits the fact that private clients consistently trade in the wrong direction and hence forecasts that account for this behavior can be highly profitable.[24] Similar behavior to a lesser extent can be attributed to CO.

To provide a visual illustration of the daily results, Figure 2 shows the out-of-sample cumulative wealth over time for the four individual order flow models relative to the RW benchmark starting with an initial wealth of $1. The figure indicates that the cumulative wealth for each of the four order flow strategies is higher than the RW. HF and PC perform better than AM and CO. More importantly, while at the beginning the order flow models tend to comove with the RW, after the crisis in 2007 the order flow models considerably outperform the RW. Hence it would be reasonable to conclude that much of the order flow prior to the crisis was 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.

The monthly results reported in Table 5 and Figure 3 are qualitatively similar to the daily results and confirm the superior performance of strategies conditioning on order flow. Note that the Sharpe ratios tend to be higher for monthly rebalancing, which is partly due to the lower transaction costs incurred as rebalancing takes place less often. The best performers among the four customer groups are AM ($\mathcal{SR} = 0.84$) and HF ($\mathcal{SR} = 0.61$). An explanation for the very good performance of AM could be, for example, that asset managers are long-term investors who may trade infrequently and hence monthly order flow is a more powerful predictor than daily order flow. It is indeed likely that informed traders choose to trade gradually, hiding their trades among the uninformed trades, in order to avoid a rapid adjustment in the price they pay.[25] Note also that the PC strategy now has by far the worse performance ($\mathcal{SR} = -0.44$) indicating that the information of private clients is very short-lived.

5.3 Order Flow and Macroeconomic Information

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.

[24] Also note that in unreported results we find that the PC strategy has a higher turnover than AM and HF. In other words, a strategy using the information of private clients trades more often and in the wrong direction.

[25] This 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.

We address this question by setting up a framework based on the set of trading strategies that condition on order flow or on standard 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). Recall 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 information models that deliver a set of out-of-sample one-period ahead forecasts for exchange rate returns. The out-of-sample forecasts are for the period of January 2004 to May 2011. In this step, it is important to note that we impose constraints on the sign of the slope parameters estimated in the predictive regressions used to generate the out-of-sample forecasts. The constraints are consistent with the economic theory that each of the eight predictive regressions represents: the slope parameters are set to be positive for UIP, TR, PPP, MF and MOM, and negative for FP and NXA. This 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. 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.

In addition to helping us understand what drives order flow, this framework will shed light on questions such as: what strategies do asset managers, hedge funds, corporates and private clients follow? Are these strategies different among customer types? Can the predictive information content in 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 behavior of FX traders and the models or information that different customers employ when deciding what assets 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., Pojarlev and Levich, 2008). 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, not the return of particular funds. Second, recent theoretical literature formalizes the notion that order flow conveys fundamental information about exchange rates and, hence, it aggregates disperse economic information (e.g., Evans and Lyons, 2007, 2008; Bacchetta and van Wincoop, 2004, 2006). This implies that order flow ought to be empirically related to macroeconomic information, and the former effectively summarizes the latter.

Table 6 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. We focus on monthly returns as this is the frequency at which most macroeconomic information is released. Consequently, the out-of-sample period of January 2004 to May 2011 comprises only 89 monthly observations. For robustness, therefore, we also report results on daily regressions in the separate Appendix. 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 for the latter three models (RW, UIP and TR), the key piece of predictive information is the interest rate differential, and hence the returns from these three 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. First, for the regression that includes all eight macroeconomic strategies, the $\overline{R}^2$ ranges from 17.9% (AM) to 48.4% (CO). The $\overline{R}^2$ for AM and HF is lower than for CO and PC indicating a lower dependence of the AM and HF trading strategies to macroeconomic information. Overall, the macroeconomic information strategies capture up to 50% of the net demand for currency manifested in the order flows. Second, the betas on the macroeconomic strategies tend to be positive but insignificant. Momentum is an exception in that all order flow strategies load negatively on it, which implies that customers often follow contrarian strategies that buy depreciating currencies and sell appreciating currencies. Notably, HF order flow loads significantly positively on UIP, and so is CO on PPP and PC on RW. Third, the alphas are positive for AM, HF and CO, negative for PC and in all cases they are insignificant. This result indicates that there is no additional excess return generated by the order flow strategies, over and above what can be generated by combining the macroeconomic information strategies.

Together these results imply that trading strategies conditioning on customer order flow can to a large extent be explained by particular combinations of macroeconomic information. Hence order flow can be partly explained but cannot be fully replicated by macroeconomic information. The absence of a significant (positive) alpha further suggests that there is no additional information in customer order flows 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.

5.3.1 Robustness to Alternative Asset Allocation Strategies

Next we evaluate the robustness of these results to the design of the asset allocation strategy by conducting four tests.[26] The first test uses naive equal weights, where we place an equal weight on all currencies for which there is a positive forecast and an equal weight on all currencies for which there is a negative forecast. The second test implements a zero investment portfolio, where the mean-variance weights are restricted to sum up to zero. The third test sets all asset correlations to be equal to zero. The fourth and final test uses a shrinkage covariance matrix using the method of Ledoit and Wolf (2004a,b). The results are reported in Table 7. We find that overall the $\overline{R}^2$ tends to improve markedly (especially for zero correlations), some of the betas tend to be more significant, but the alphas remain insignificant. Hence our results are qualitatively the same as before.

5.3.2 Combined Forecasts Based on Macroeconomic Information

Our final exercise based on simple OLS regressions involves assessing the predictive ability of forecast combinations by optimally combining the exchange rate return forecasts of the eight macroeconomic information strategies. The rationale of this exercise is as follows. We so far find that order flow has substantial predictive ability and some of it can be explained ex post by particular combinations of macroeconomic fundamentals. It is therefore natural to ask whether a forecast combination of macroeconomic information constructed ex ante and based on well-established methods in the literature can match the predictive ability of order flow. If so, then order flow simply replicates widely used ex ante forecast combinations and does not contribute any further predictability. If not, however, 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.

Our analysis is comprehensive as it employs six forecast combinations: 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 in one predictive regression, and Principal Component Analysis (PCA) that uses the first principal component. The results, reported in Table 8, provide overwhelming evidence that out of sample almost none of the forecast combination methods can outperform the RW benchmark. This is true for both daily and monthly rebalancing. With the single exception of the daily PCA combination, all other cases display negative Sharpe ratios and highly negative performance fees. This is in sharp contrast to the predictability of order flow reported

[26] These robustness tests are similar to the ones implemented by DeMiguel, Plyakha, Uppal and Vilkov (2012).

in Tables 4 and 5. 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.

5.3.3 Time-Varying Parameter Regressions

While the OLS regressions capturing the relation between order flow and macroeconomic information are estimated with constant parameters, it is likely that these parameters are changing 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 a well-documented practice in currency markets 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, 2006), where every day 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 every day as the market rationally searches for an explanation for the observed exchange rate change.[27]

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 on monthly portfolio excess returns by conditioning out-of-sample on the eight macroeconomic strategies. 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 for small samples as described in the separate Appendix.

We begin with Figure 4, which plots the time-varying alphas across 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 story emerges as we can identify when alphas are positive, when they are negative and when they are significant. For example, after the collapse of Lehman's in September 2008, the AM and HF order flow strategies delivered a significantly positive alpha, thus performing better than a combination of macroeconomic fundamentals. In contrast, over the same period the PC strategy delivered a significantly negative alpha. Clearly, more sophisticated investors seem to have significantly outperformed the less sophisticated investors during the crisis.

The next four figures (Figs. 5 to 8) display the time-varying betas and $\overline{R}^2$ for each customer group. 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. The

[27] This practice is documented, for example, in the survey evidence of Cheung and Chinn (2001) that is based on questionnaires sent to US FX traders.

graphs in essence show the anatomy of the month-by-month trading strategies of the four customer groups. Note also that the $\overline{R}^2$ tends to be high most of the time for all customer groups, typically in the range of 80%, but exhibits negative spikes. Still, the average $\overline{R}^2$ is much higher in the TVP regressions than in the constant parameter regressions: 50% for AM, 60% for HF, and 70% for CO and PC. This is further evidence of the strong relation between order flow and macroeconomic information.

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 9 to 12 plot the cumulative monthly order flow of a given customer group dedicated to a particular macroeconomic currency investment strategy. This cumulative order flow is defined as 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 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 the macroeconomic strategies. Furthermore, a noteworthy example is that asset managers and hedge funds reduced their exposure to the carry trade at least a few months before the crisis started. These figures focus on the monthly results, whereas the daily results are shown in the separate Appendix.

5.4 Summary of Results

The empirical evidence supports the following three main findings. First, our results indicate that customer order flow has substantial predictive ability for exchange rate returns. This predictability can lead profitable trading strategies net of transaction costs. At the monthly horizon, for instance, the predictive ability of customer order flow tends to be higher for asset managers and hedge funds and lower or non-existent for corporates and private clients. This finding justifies the use of disaggregate data on customer order flow.

Second, simple 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" in generating currency orders and informing customers' trading decisions. Notably, there is no evidence that order flow has additional predictive information over and above the information embedded in macroeconomic fundamentals. In other words, there is no "alpha" in customer order flow. Furthermore, standard forecast combinations of macroeconomic variables fail to replicate *ex ante* the predictive ability of order flow. This leads us to conclude that order flow provides a distinct and effective way of aggregating macroeconomic information with substantial ex ante predictive ability.

Third, the relation between order flow and macroeconomic information can vary significantly over time as investors react to changes in their macroeconomic environment by assigning a different weight each month to different macroeconomic variables. Time-varying parameter regressions provide a more complete picture of how customer order flow relates to different macroeconomic fundamentals. The results are consistent with the scapegoat theory of Bacchetta and van Wincoop (2004, 2006). In these time-varying parameter regressions, macroeconomic information can explain 50% to 70% of the predictive ability of order flow.

6 Conclusions

Trades between customers and FX dealers generate a measure of order flow that conveys the customers' information on exchange rates. Dealers can then act on this information and reveal it to the rest of the FX market through interdealer trading. This mechanism implies that customer order flow may be able to predict future exchange rate returns. In this paper, we examine whether this is the case using a unique data set on daily order flow representing the transactions of customers and UBS, a top FX dealer globally. The data set ranges from 2001 to 2011, covers the G10 currencies, more importantly, is disaggregated across four different end-user segments of the FX market: asset managers, hedge funds, corporate clients, and private clients.

The empirical analysis first assesses the predictive ability of customer order flow on FX excess returns in the context of dynamic asset allocation strategies. We then relate the portfolio returns generated from order flow strategies to the portfolio returns of other strategies based on public macroeconomic information such as interest rates, real exchange rates or monetary fundamentals.

We find that the dynamic trading strategy based on customer order flow strongly outperforms the popular carry trade in sample and out of sample. More importantly, the portfolio returns generated from conditioning on order flow can be largely explained using a combination of strategies based on publicly available macroeconomic information. This is especially true when we allow for time-variation in the relation of order flow to macroeconomic fundamentals. Furthermore, there is no "alpha" in customer order flow as it contains no additional predictive information over and above the macroeconomic information. These findings support the notion that order flow aggregates disperse public information about economic fundamentals that are relevant to exchange rates.

Table 1: The FX Market Share of UBS

The table displays the overall market share and the market share by customer type for UBS bank. 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) are presented for *real money, leveraged funds* and *non-financial corporations*.

	Overall Market		Real Money		Leveraged Funds		Non-financial 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. Daily 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.

		Mean	*Sdev*	*Min*	*Max*	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. Daily 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 Daily Order Flow

The table reports the in-sample and out-of-sample economic value of the predictive ability of daily order flow. The results are based on dynamic asset allocation strategies investing in the G-10 currencies with daily rebalancing. The benchmark strategy is the naïve random walk (RW) model. The competing 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 transaction costs. The strategy invests in a domestic bond and nine foreign bonds and is rebalanced daily. For each strategy, we report the annualized mean (r_p), annualized volatility (σ_p), skewness (*Skew*), excess kurtosis (*Kurt*), annualized Sharpe ratio (*SR*), annualized Sortino ratio (*SO*), maximum drawdown (*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.

Strategy	r_p (%)	σ_p (%)	*Skew*	*Kurt*	ρ_1	$\mathcal{SR}$	$\mathcal{SO}$	$\mathcal{MDD}$ (%)	$\mathcal{P}$ (*bps*)
				In-Sample Period: Jan 2001 - May 2011					
RW	2.9	9.6	−0.70	11.33	−0.06	0.13	0.15	37.4	
AM	8.6	10.7	−0.17	5.37	−0.07	0.64	0.89	19.4	503
HF	10.9	9.9	0.05	2.26	−0.01	0.94	1.43	14.5	795
CO	10.5	10.3	−0.15	2.00	−0.02	0.86	1.28	16.9	728
PC	9.2	10.2	−0.24	2.56	−0.02	0.74	1.08	22.9	603
TOT	9.2	10.3	−0.09	3.35	−0.04	0.73	1.05	19.4	589
ALL	12.0	10.5	−0.17	2.80	−0.06	0.98	1.42	14.2	860
				Out-of-Sample Period: Jan 2004 - May 2011					
RW	−0.8	14.4	−0.70	7.26	−0.05	−0.17	−0.21	47.0	
AM	1.3	13.2	−0.47	2.47	0.02	−0.03	−0.04	37.6	315
HF	4.9	12.2	−0.46	2.06	0.04	0.27	0.36	36.2	759
CO	4.0	13.7	−0.60	2.69	−0.01	0.17	0.22	39.2	539
PC	6.9	13.4	−0.58	2.50	−0.01	0.40	0.52	35.2	868
TOT	4.5	12.7	−0.52	2.47	0.03	0.22	0.29	34.7	676
ALL	4.4	12.1	−0.51	2.77	0.01	0.23	0.31	37.7	713

Table 5. The Predictive Ability of Monthly 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 G-10 currencies with monthly rebalancing. The benchmark strategy is the naïve random walk (RW) model. The competing 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 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 (*SR*), annualized Sortino ratio (*SO*), maximum drawdown (*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.

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

Table 6. Monthly 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 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. The regressions shown below impose constraints on the sign of the parameters of the predictive regressions used to generate the forecasts. The sign constraints are consistent with economic theory: all parameters 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}	$\overline{R}^2$
					Out-of-Sample Period: Jan 2004 - May 2011					
AM	0.078	0.295				0.599	−0.116	−0.123	−0.100	16.5
	(0.052)	(0.141)				(0.280)	(0.113)	(0.173)	(0.103)	
	0.093			−0.231		0.834^{b}	−0.075	−0.056	−0.134	12.2
	(0.054)			(0.157)		(0.245)	(0.122)	(0.167)	(0.122)	
	0.086				0.168	0.554	−0.106	−0.024	−0.159	8.6
	(0.057)				(1.943)	(1.692)	(0.118)	(0.164)	(0.119)	
	0.094	0.124	0.224	−0.228	0.099	0.522	−0.078	−0.191	−0.106	17.9
	(0.053)	(0.182)	(0.174)	(0.138)	(1.886)	(1.606)	(0.088)	(0.162)	(0.113)	
HF	0.045	0.276				0.633^{b}	0.038	−0.150	−0.020	15.4
	(0.051)	(0.139)				(0.241)	(0.120)	(0.101)	(0.099)	
	0.047			0.278		0.579^{b}	−0.002	−0.011	−0.098	13.9
	(0.053)			(0.141)		(0.237)	(0.124)	(0.084)	(0.104)	
	0.047				0.745	0.059	0.068	−0.072	−0.088	9.8
	(0.060)				(1.031)	(0.872)	(0.124)	(0.095)	(0.127)	
	0.022	0.323	0.014	0.376^{a}	1.237	−0.724	0.021	−0.143	−0.073	22.3
	(0.049)	(0.153)	(0.192)	(0.160)	(1.317)	(1.259)	(0.132)	(0.127)	(0.104)	
CO	0.032	0.352^{b}				1.240^{c}	−0.095	−0.113	−0.240	40.2
	(0.049)	(0.119)				(0.188)	(0.104)	(0.169)	(0.139)	
	0.037			0.223		1.245^{c}	−0.125	0.044	−0.327	34.7
	(0.053)			(0.173)		(0.224)	(0.131)	(0.186)	(0.162)	
	0.069				-2.513^{a}	3.643^{b}	−0.178	0.069	−0.251	39.9
	(0.051)				(0.947)	(0.857)	(0.131)	(0.166)	(0.141)	
	0.053	0.311	0.110	0.205	−2.144	2.995^{a}	−0.204	−0.061	−0.216	48.4
	(0.045)	(0.150)	(0.170)	(0.137)	(1.02)	(0.969)	(0.114)	(0.153)	(0.124)	
PC	−0.102	0.738^{a}				0.203	−0.155	0.120	−0.158	39.9
	(0.066)	(0.204)				(0.314)	(0.188)	(0.198)	(0.129)	
	−0.081			0.137		0.395	−0.166	0.398	−0.310	13.1
	(0.079)			(0.226)		(0.374)	(0.285)	(0.304)	(0.197)	
	−0.056				−2.108	2.379	−0.219	0.427	−0.251	16.6
	(0.065)				(2.286)	(1.936)	(0.316)	(0.309)	(0.174)	
	−0.085	0.697^{a}	0.111	0.207	−1.707	1.560	−0.249	0.162	−0.144	44.1
	(0.052)	(0.234)	(0.199)	(0.183)	(1.970)	(1.646)	(0.242)	(0.211)	(0.115)	

Table 7. Monthly Order Flow and Macroeconomic Information: Robustness to Alternative Asset Allocation Strategies

This table shows the results of four robustness tests in assessing 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 regressions are set up exactly as described in Table 6. *Naive equal weights* is the case of placing an equal weight on all currencies for which there a positive forecast and a (potentially different) equal weight on all currencies for which there a negative forecast. *Zero investment portfolio* is the mean-variance case where the weights are restricted to sum up to zero. *Zero correlations* is the mean-variance case where all correlations between the risky assets are set to be equal to zero. *Shrinkage covariance matrix* is the mean-variance case where the covariance matrix is optimally shrunk using the method of Ledoit and Wolf (2004a,b). The out-of-sample period used to generate portfolio returns runs from January 2004 to May 2011.

	α	β_{RW}	β_{FP}	β_{UIP}	β_{TR}	β_{PPP}	β_{MF}	β_{NXA}	β_{MOM}	$\overline{R}^2$
					Out-of-Sample Period: Jan 2004 - May 2011					
					Naive Equal Weights					
AM	−0.010	0.427^a	0.371^a	0.135	0.491	−0.457	−0.493	0.196	−0.111	60.1
	(0.026)	(0.144)	(0.150)	(0.125)	(0.455)	(0.463)	(0.188)	(0.115)	(0.134)	
HF	−0.014	0.089	0.468^b	0.190	−0.176	0.343	0.061	0.114	0.128	46.2
	(0.027)	(0.154)	(0.147)	(0.133)	(0.461)	(0.427)	(0.169)	(0.155)	(0.107)	
CO	0.067^a	0.389	0.233	−0.083	−0.017	0.107	0.062	−0.158	−0.192	31.3
	(0.031)	(0.183)	(0.143)	(0.144)	(0.527)	(0.522)	(0.186)	(0.160)	(0.129)	
PC	−0.009	0.283	0.472^b	0.110	−0.275	0.272	−0.299	0.212	−0.012	57.1
	(0.025)	(0.190)	(0.142)	(0.117)	(0.402)	(0.392)	(0.232)	(0.163)	(0.105)	
					Zero Investment Portfolio					
AM	0.061	0.235	0.097	−0.243	0.188	−0.028	−0.015	−0.271	−0.237	26.4
	(0.042)	(0.147)	(0.156)	(0.125)	(1.513)	(1.383)	(0.126)	(0.159)	(0.138)	
HF	0.012	0.273	0.097	0.213	0.760	−0.477	0.073	−0.141	−0.134	10.7
	(0.042)	(0.193)	(0.212)	(0.179)	(0.992)	(0.991)	(0.153)	(0.140)	(0.119)	
CO	0.035	0.474^b	−0.039	0.186	-2.091^b	2.619^c	−0.208	−0.080	−0.169	44.1
	(0.040)	(0.142)	(0.147)	(0.115)	(0.735)	(0.696)	(0.109)	(0.150)	(0.109)	
PC	-0.107^b	0.833^c	−0.092	−0.049	−1.949	1.258	−0.194	0.004	0.063	53.6
	(0.041)	(0.182)	(0.169)	(0.166)	(1.376)	(1.296)	(0.191)	(0.157)	(0.119)	
					Zero Correlations					
AM	0.053	0.220	0.255	0.029	0.371	0.248	-0.292^a	0.322^a	−0.025	84.3
	(0.036)	(0.121)	(0.126)	(0.104)	(1.310)	(1.262)	(0.115)	(0.129)	(0.084)	
HF	0.009	0.581^b	−0.233	0.099	0.041	0.808	−0.148	0.069	−0.024	82.6
	(0.039)	(0.167)	(0.155)	(0.136)	(1.221)	(1.215)	(0.141)	(0.139)	(0.083)	
CO	0.056	0.041	0.275	−0.162	−0.275	1.338	0.000	0.080	−0.169	84.2
	(0.035)	(0.200)	(0.220)	(0.102)	(1.077)	(1.085)	(0.130)	(0.151)	(0.084)	
PC	0.010	0.335	0.082	0.096	0.488	0.332	−0.243	0.303	−0.097	89.7
	(0.031)	(0.203)	(0.198)	(0.111)	(1.071)	(1.060)	(0.149)	(0.135)	(0.076)	
					Shrinkage Covariance Matrix					
AM	0.086^a	0.209	0.199	−0.139	−0.411	0.975	−0.161	−0.069	−0.087	21.8
	(0.042)	(0.166)	(0.168)	(0.143)	(1.671)	(1.522)	(0.130)	(0.142)	(0.122)	
HF	0.027	0.380	−0.009	0.414^a	0.885	−0.400	−0.012	−0.158	−0.072	24.6
	(0.042)	(0.182)	(0.192)	(0.172)	(1.205)	(1.135)	(0.136)	(0.135)	(0.105)	
CO	0.055	0.336^a	0.120	0.147	−2.244	3.048^a	−0.207	0.002	−0.226	48.0
	(0.040)	(0.152)	(0.158)	(0.121)	(1.067)	(1.039)	(0.129)	(0.149)	(0.116)	
PC	−0.080	0.685^b	0.057	0.22	−0.464	0.611	−0.257	0.244	−0.197	49.9
	(0.042)	(0.213)	(0.173)	(0.157)	(1.717)	(1.539)	(0.244)	(0.181)	(0.117)	

Table 8. The Predictive Ability of Combined Forecasts Conditioning on Macroeconomic Information

The table reports the daily and monthly 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 G-10 currencies with daily or 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 transaction costs. The strategy invests in a domestic bond and nine foreign bonds and is rebalanced daily or monthly. For each strategy, we report the annualized mean (r_p), annualized volatility (σ_p), skewness (*Skew*), excess kurtosis (*Kurt*), annualized Sharpe ratio (*SR*), annualized Sortino ratio (*SO*), maximum drawdown (*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.

Strategy	r_p (%)	σ_p (%)	*Skew*	*Kurt*	*AR*(1)	$\mathcal{SR}$	$\mathcal{SO}$	$\mathcal{MDD}$ (%)	*TO* (%)	$\mathcal{P}$ (*bps*)
					Out-of-Sample Period: Jan 2004 - May 2011					
					Daily Results					
RW	−0.8	14.4	−0.70	7.26	−0.05	−0.17	−0.21	−47.04	2.5	
AVE	−6.5	18.6	−0.46	14.76	−0.08	−0.44	−0.57	−49.84	30.3	−1000
MED	−2.2	19.4	−0.46	3.59	0.01	−0.20	−0.26	−51.16	19.3	−651
TRI	−2.4	15.4	−0.13	1.61	0.01	−0.26	−0.38	−48.31	26.0	−241
MSE	−8.3	20.3	−0.28	14.14	−0.07	−0.49	−0.63	−55.39	27.9	−1383
KS	−4.6	19.1	−0.38	13.40	−0.08	−0.33	−0.43	−41.24	22.8	−862
PCA	2.4	15.5	−0.27	2.70	0.01	0.05	0.06	−34.73	14.3	230
					Monthly Results					
RW	5.3	13.5	−0.28	0.55	0.19	0.21	0.32	−40.05	0.32	
AVE	−2.7	19.6	−0.44	5.12	−0.30	−0.27	−0.36	−36.26	1.98	−1585
MED	1.4	20.3	−0.32	4.17	−0.18	−0.06	−0.08	−42.78	1.90	−1240
TRI	−2.8	18.2	−0.39	2.46	−0.29	−0.29	−0.42	−40.56	1.61	−1336
MSE	−2.1	22.1	−0.88	5.20	−0.16	−0.21	−0.26	−54.42	2.37	−2068
KS	−0.8	20.2	−0.22	3.45	−0.24	−0.17	−0.23	−34.20	2.12	−1412
PCA	−0.9	19.4	−0.65	1.58	−0.23	−0.18	−0.22	−46.66	1.43	−1316

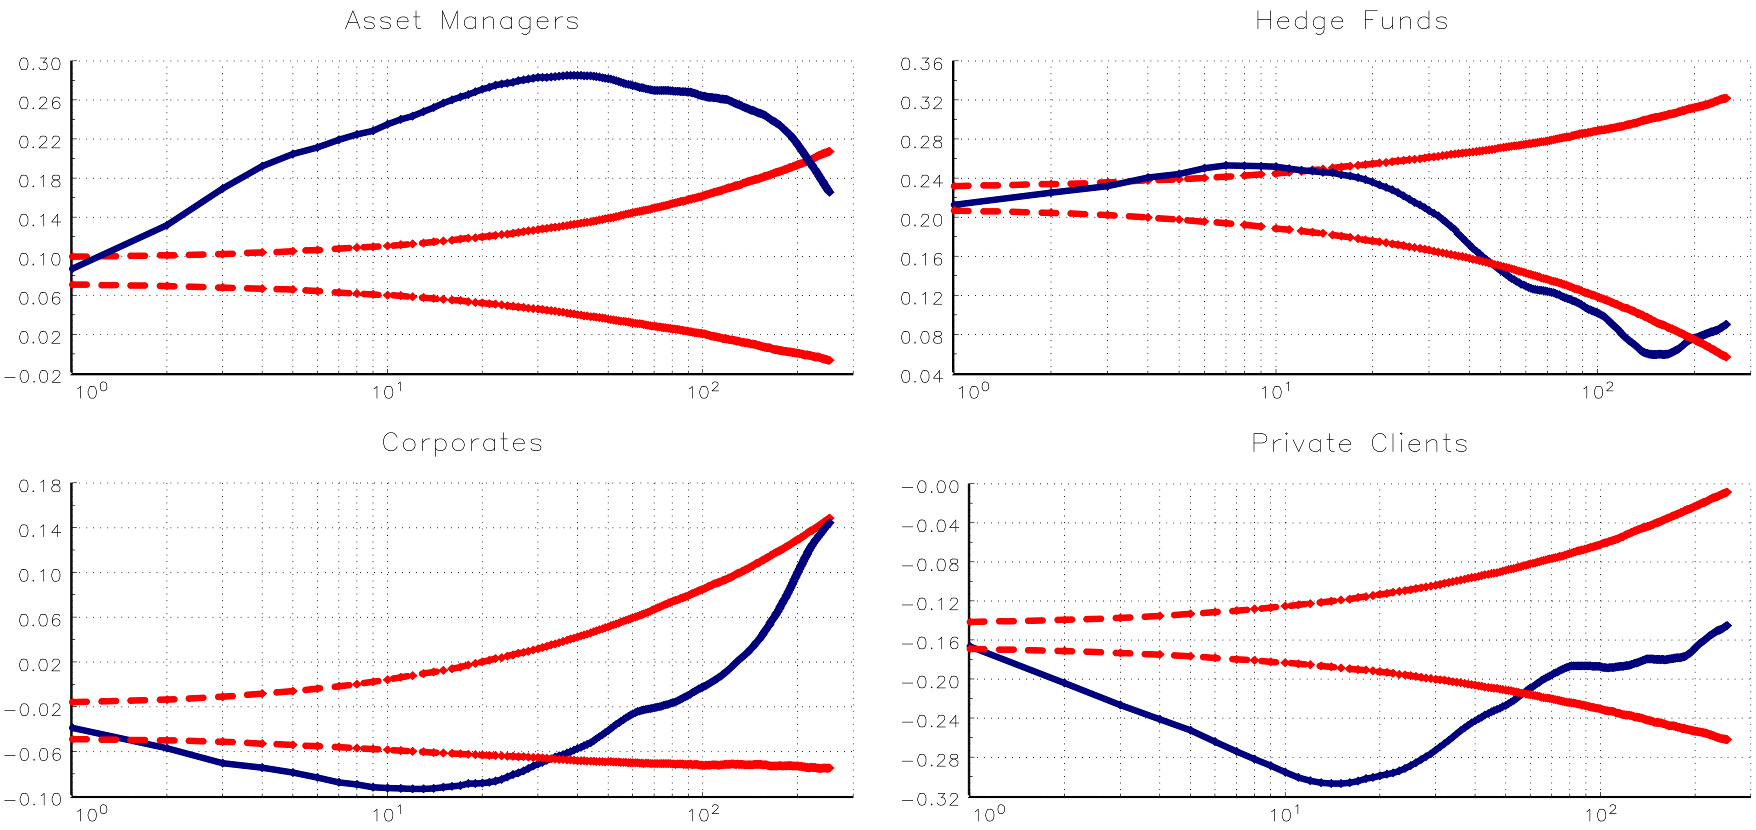

Figure 1. Correlation between FX Excess Returns and Currency Order Flow across Horizons

This figure plots the contemporaneous correlation coefficient (vertical axis) between exchange rate returns and customer order flow 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 90th percentile bootstrap confidence intervals, estimated by generating 10,000 replications. The exchange rate is defined as the US dollar price of a unit of foreign currency so that an increase (decrease) in the exchange rate implies an appreciation (depreciation) of the foreign currency. The 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 flow is classified into four customer segments: *asset managers*, *hedge funds*, *corporates* and *private clients*. The sample period ranges from from January 2001 to May 2011. Exchange rates are from Thomson Reuters while customer order flows are proprietary data from UBS.

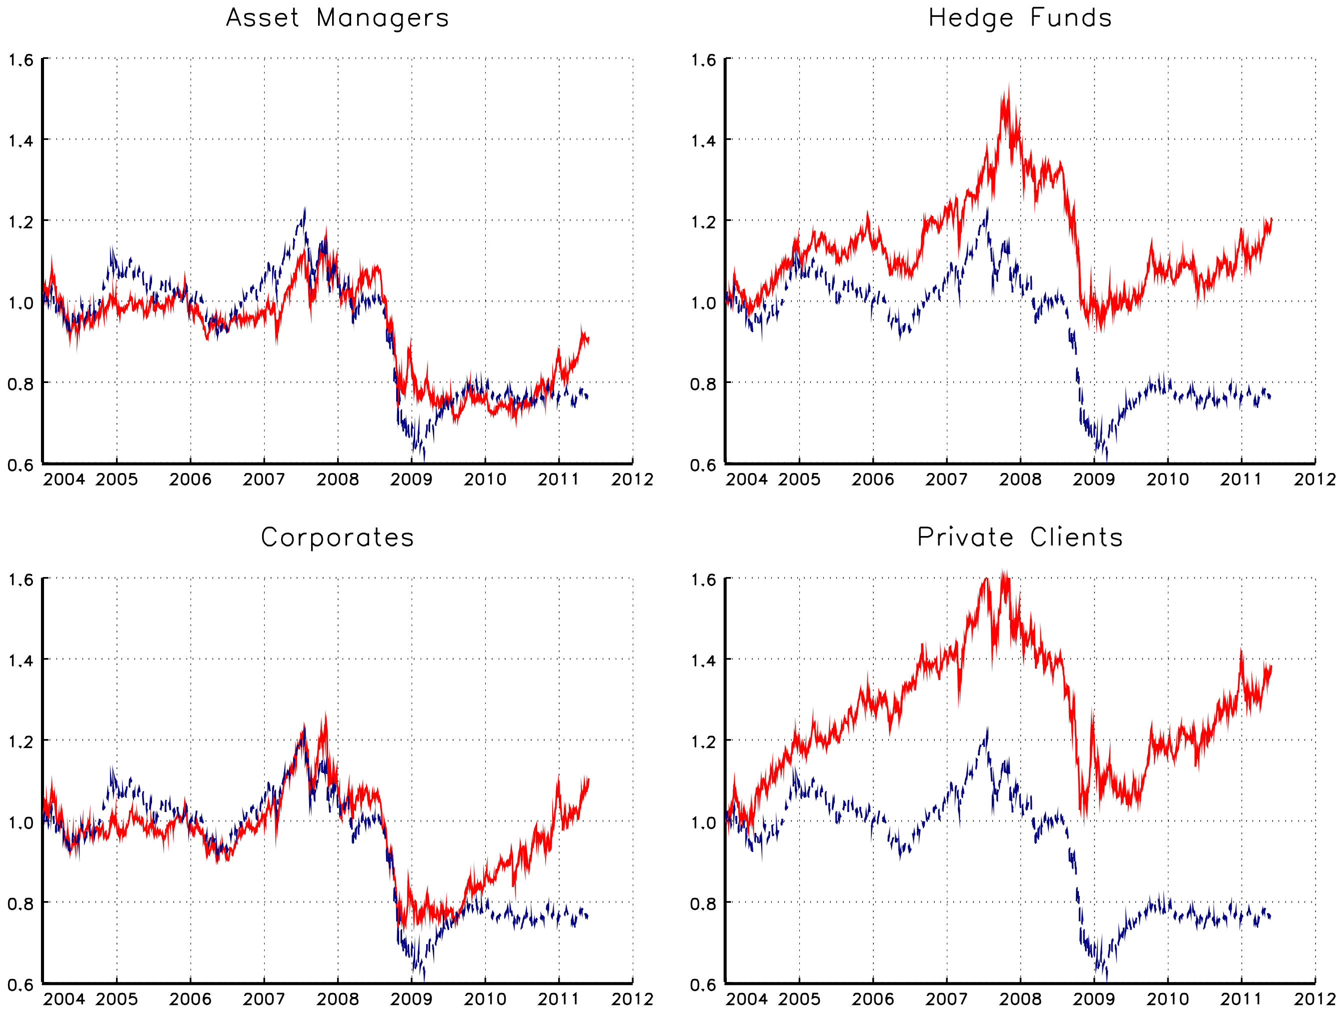

Figure 2. Cumulative Wealth of Daily Currency Order Flow Strategies

This figure displays the out-of-sample cumulative wealth of dynamic investment strategies conditioning on daily currency order flow (solid red line) relative to the naïve random walk benchmark (dashed blue line). Currency order flow is defined as the difference between the value of buyer-initiated and seller-initiated transactions and are 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.

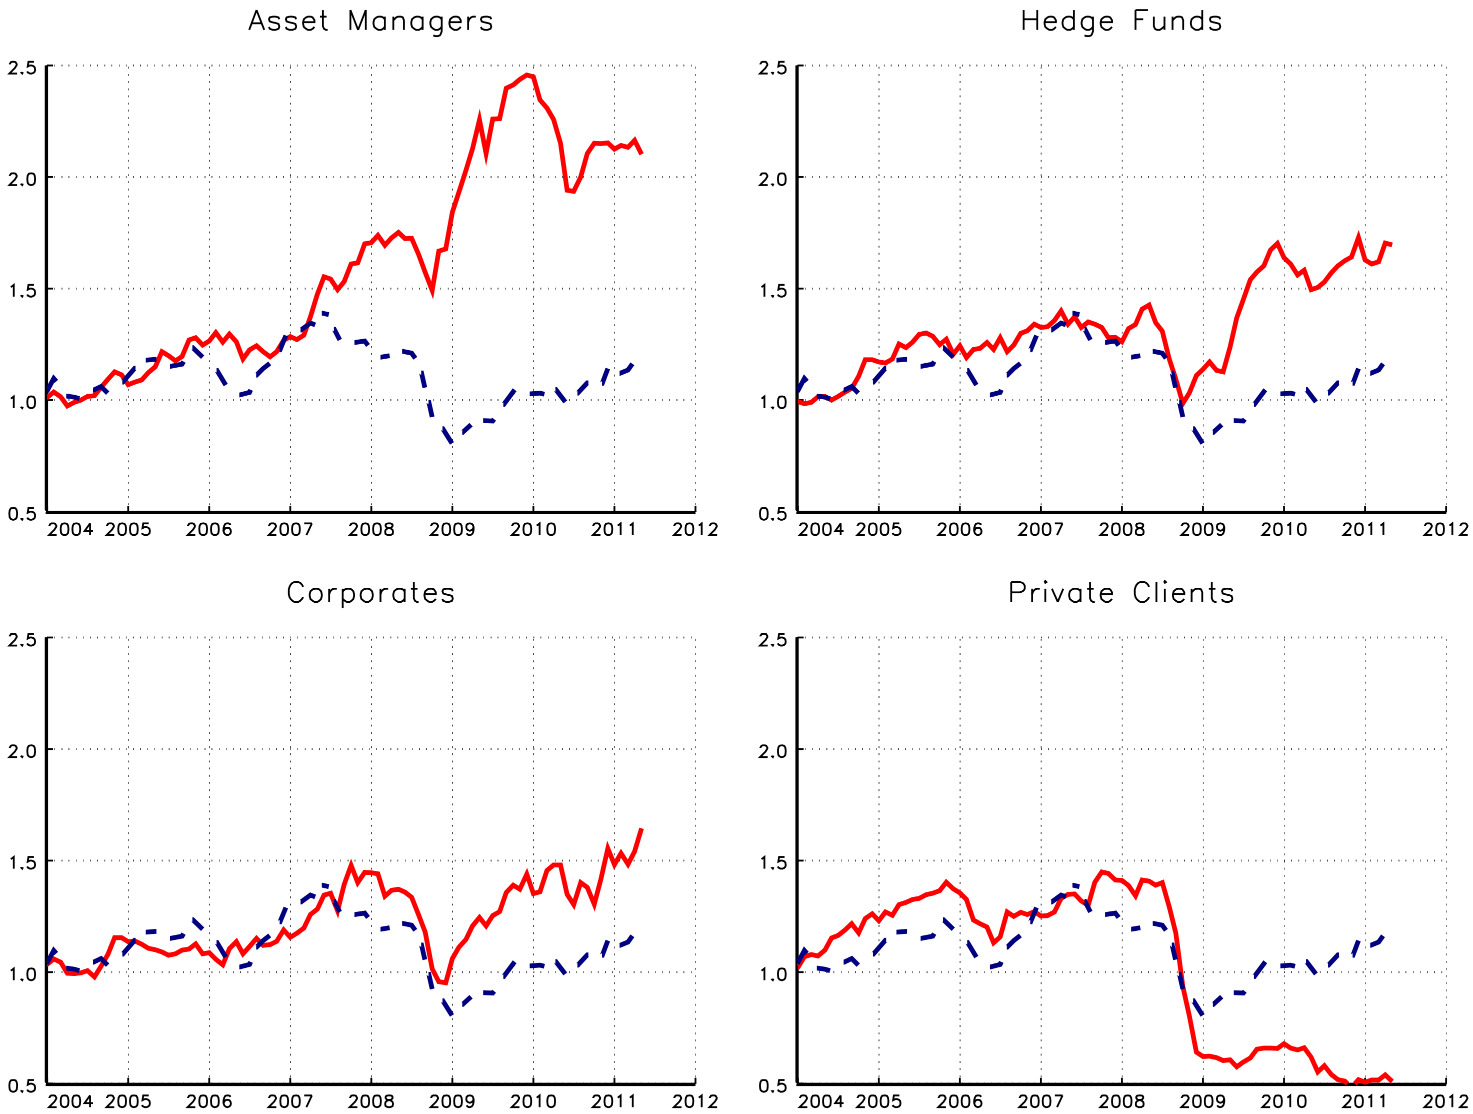

Figure 3. Cumulative Wealth of Monthly Currency Order Flow Strategies

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). Currency order flow is defined as the difference between the value of buyer-initiated and seller-initiated transactions and are 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.

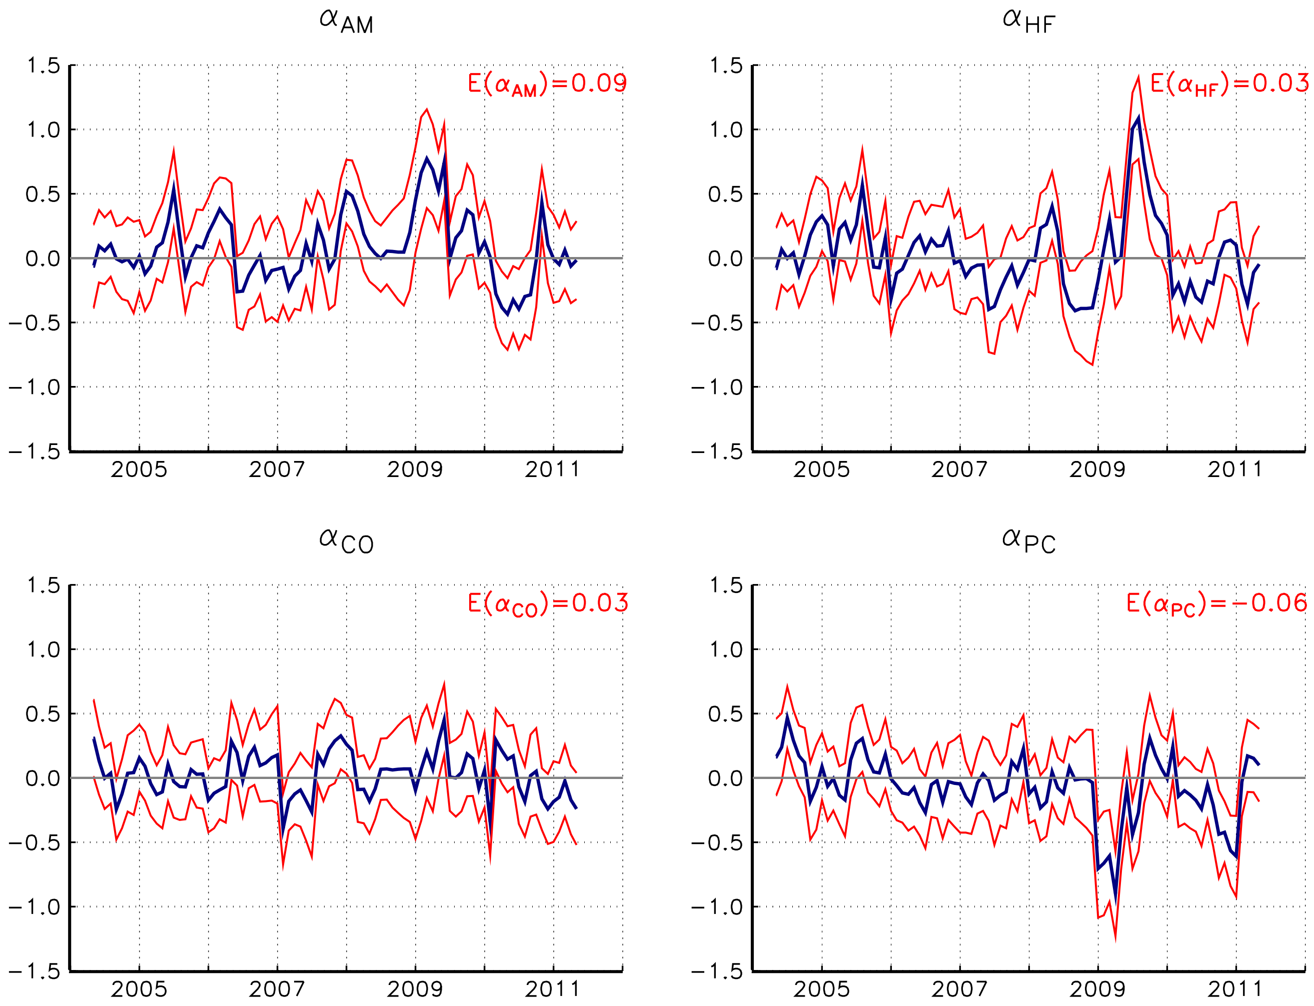

Figure 4. Monthly Order Flow and Macroeconomic Information: Time-Varying Alphas

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.

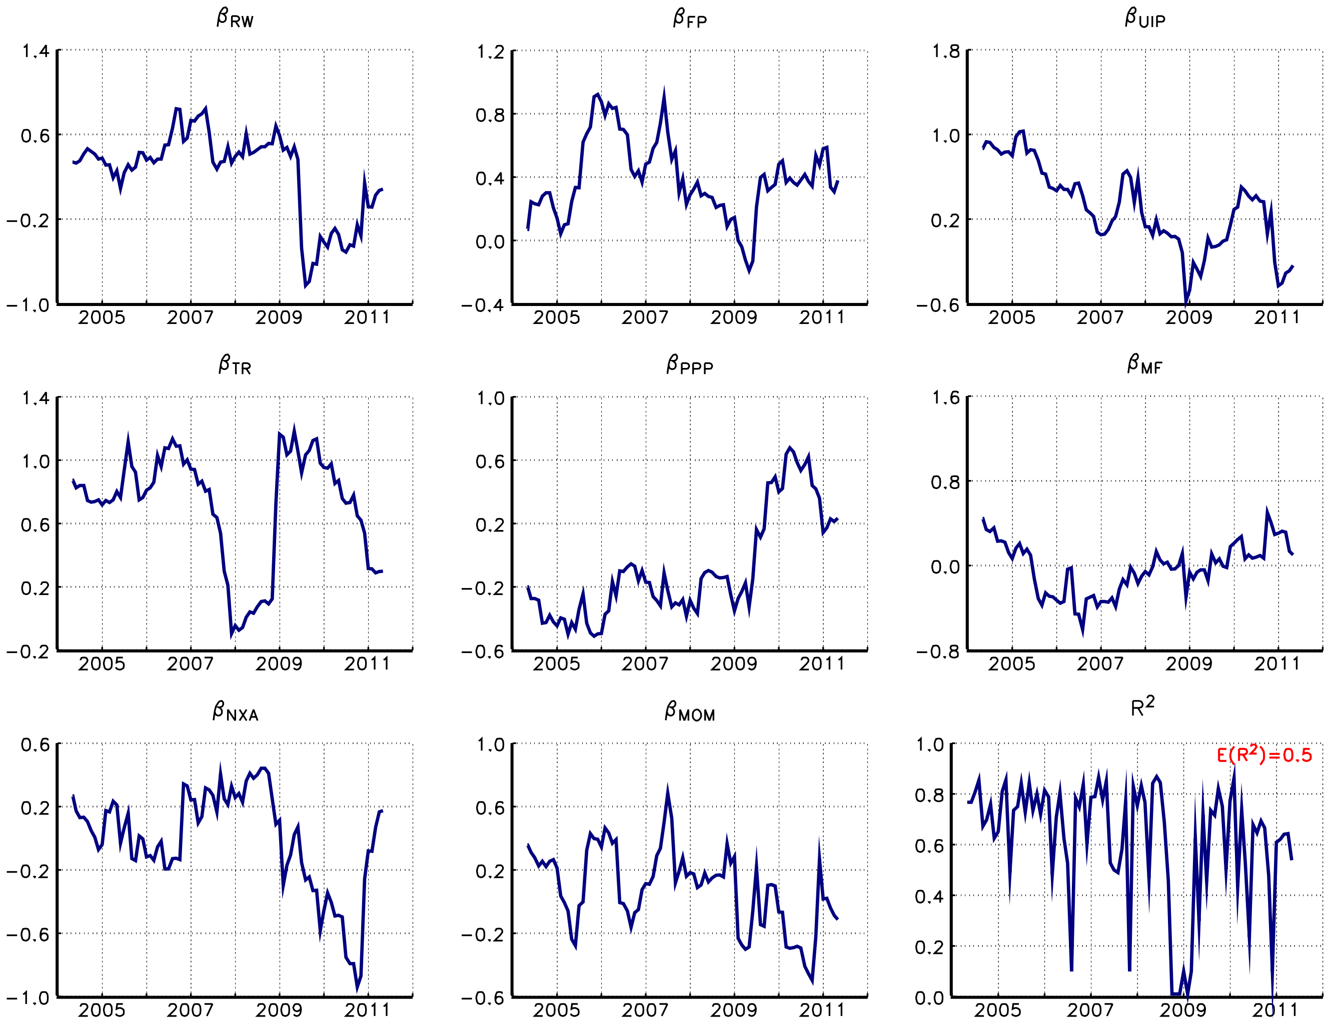

Figure 5. Time-Varying Betas: Asset Managers

The figure displays the time-varying betas (β) and adjusted R^2 on the relation between monthly currency order flow and macroeconomic information for asset managers. The betas and adjusted 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.

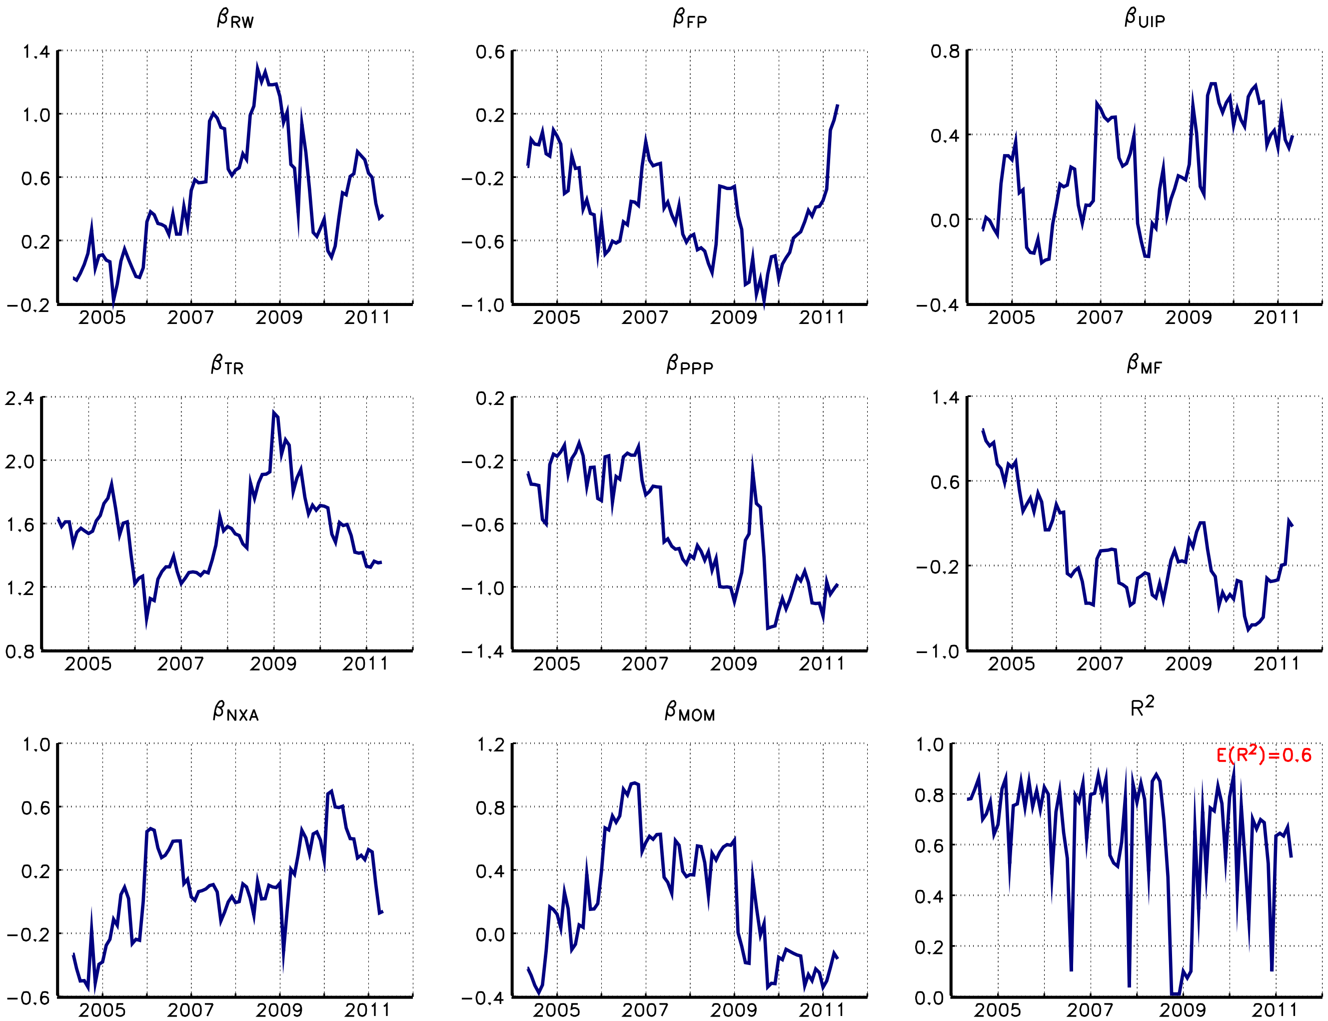

Figure 6. Time-Varying Betas: Hedge Funds

The figure displays the time-varying betas (β) and adjusted R^2 on the relation between monthly currency order flow and macroeconomic information for hedge funds. The betas and adjusted 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.

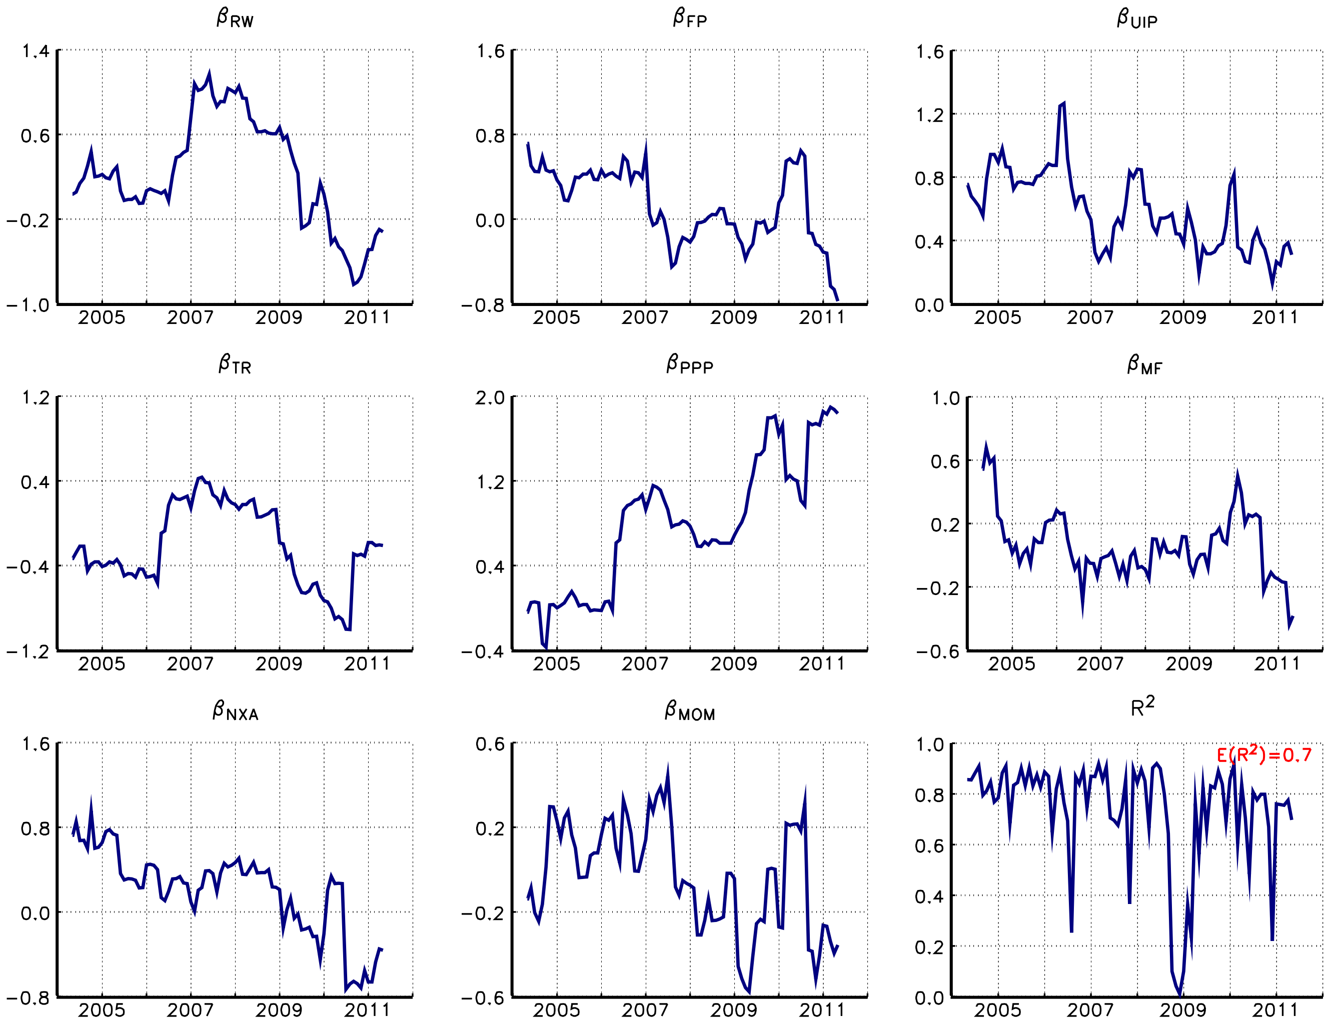

Figure 7. Time-Varying Betas: Corporates

The figure displays the time-varying betas (β) and adjusted R^2 on the relation between monthly currency order flow and macroeconomic information for corporates. The betas and adjusted 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.

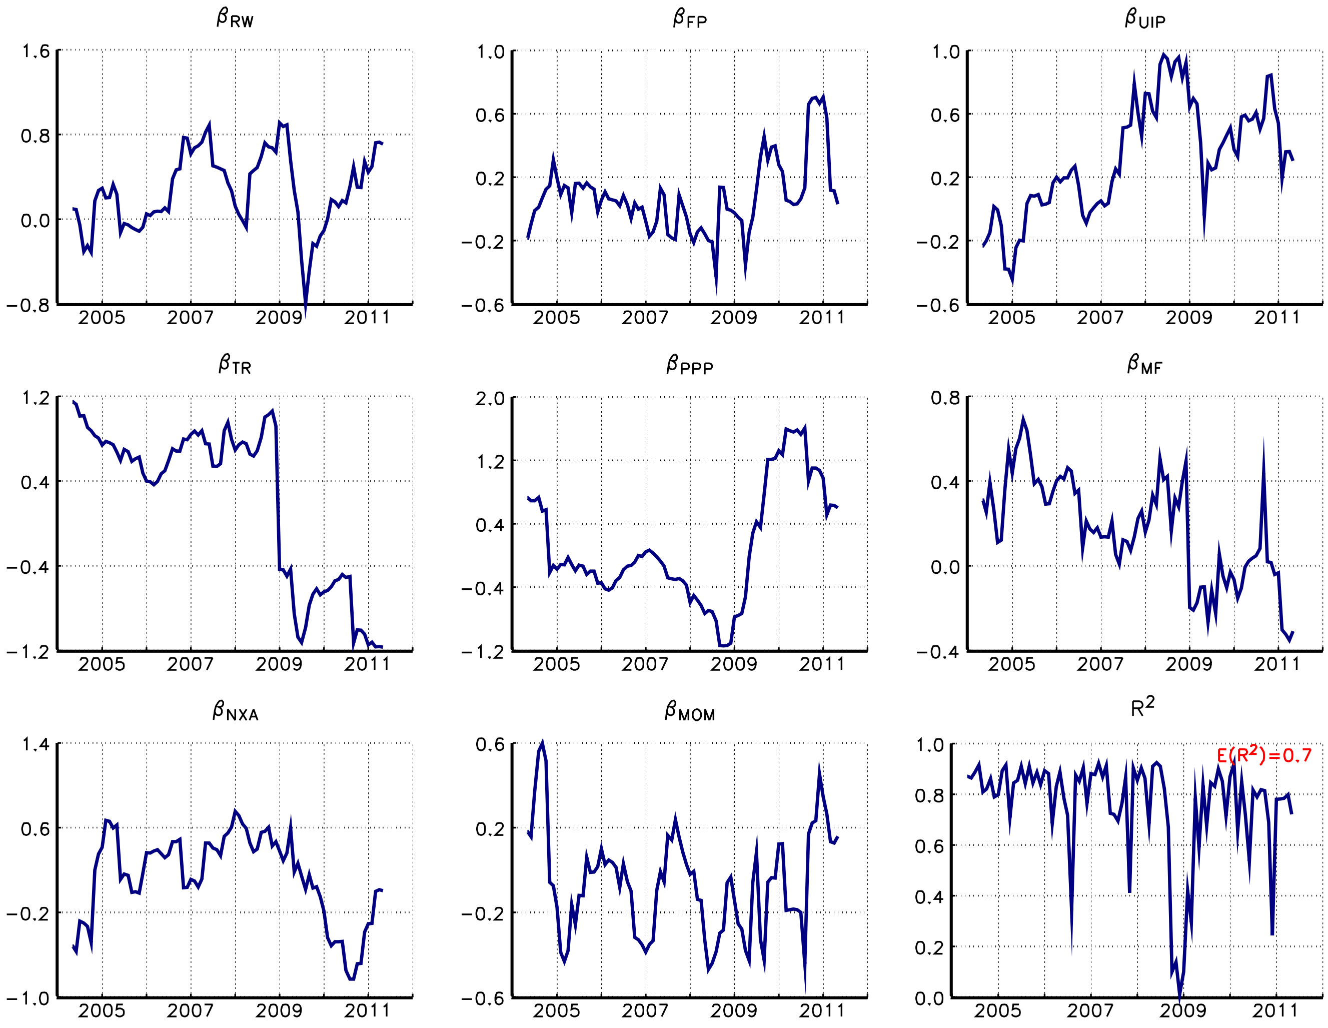

Figure 8. Time-Varying Betas: Private Clients

The figure displays the time-varying betas (β) and adjusted R^2 on the relation between monthly currency order flow and macroeconomic information for private clients. The betas and adjusted 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.

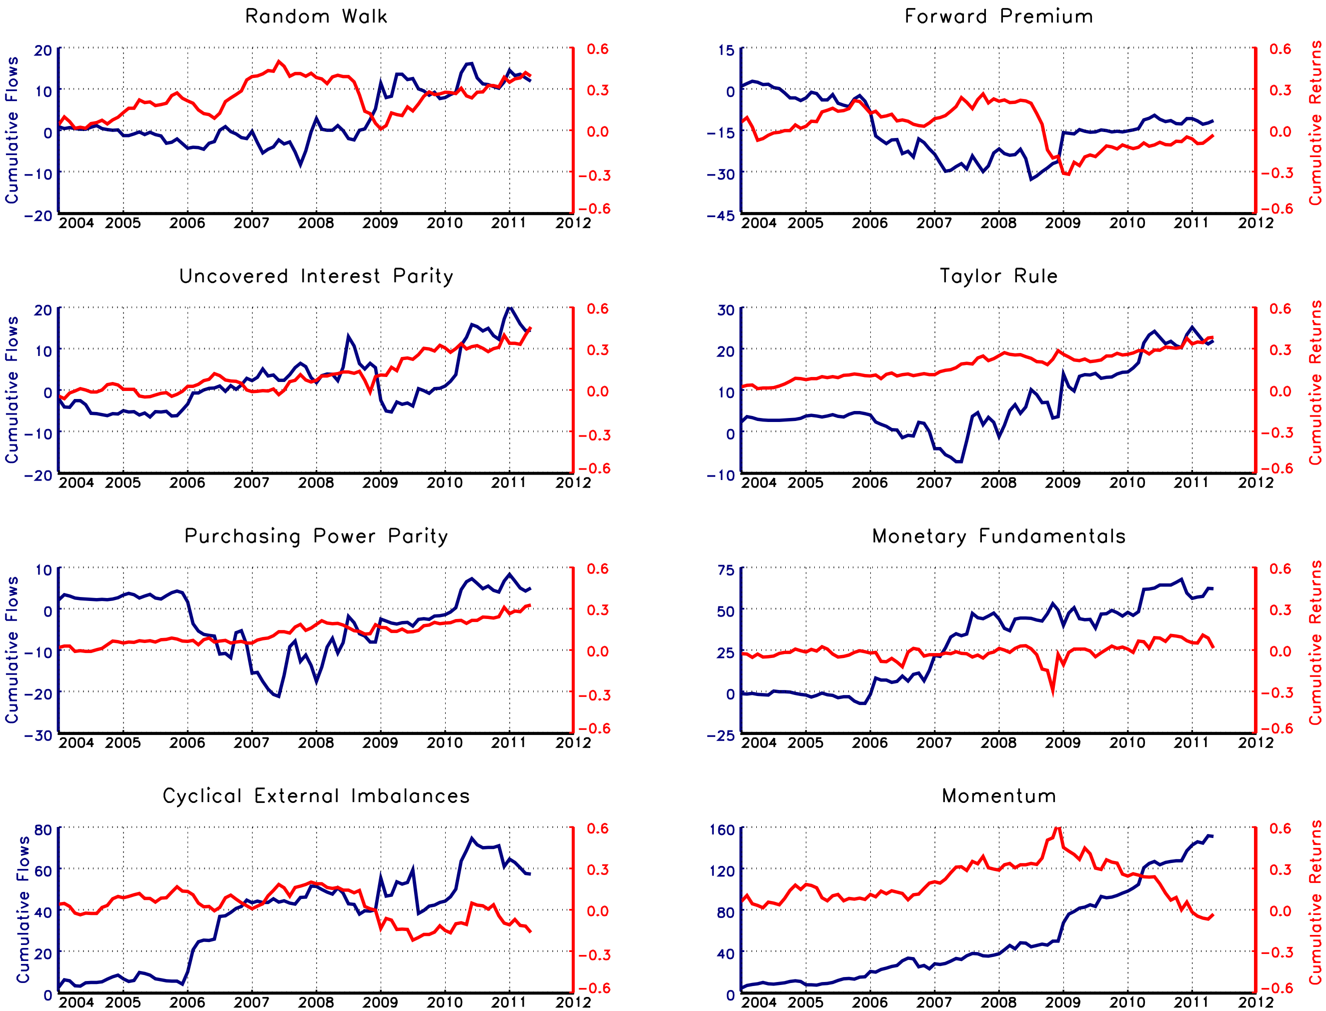

Figure 9. Cumulative Monthly Order Flow for Macroeconomic Strategies: Asset Managers

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 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.

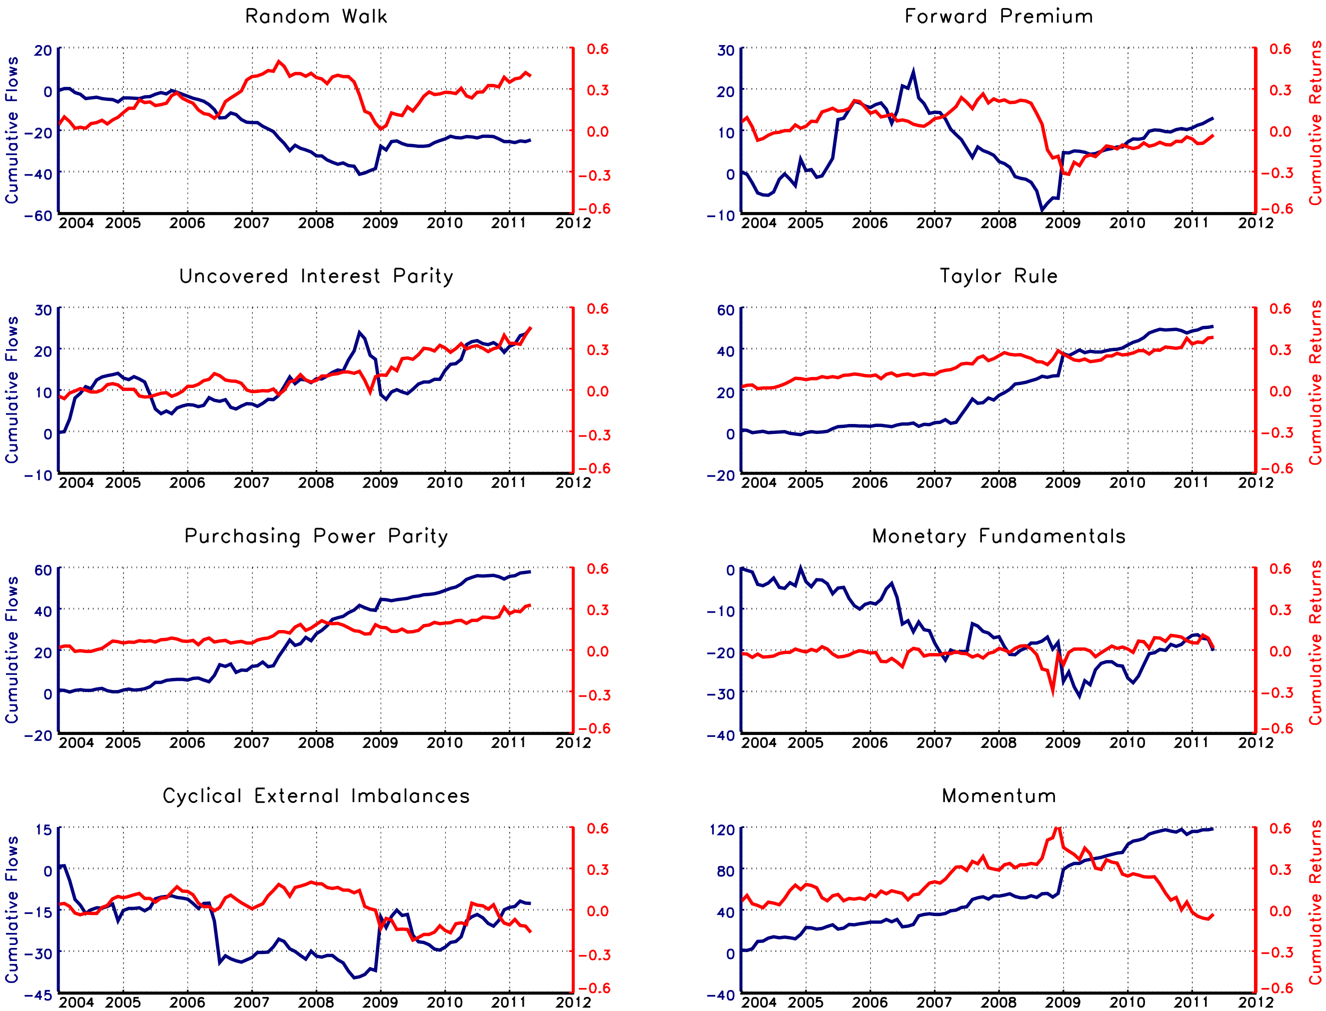

Figure 10. Cumulative Monthly Order Flow for Macroeconomic Strategies: Hedge Funds

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 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.

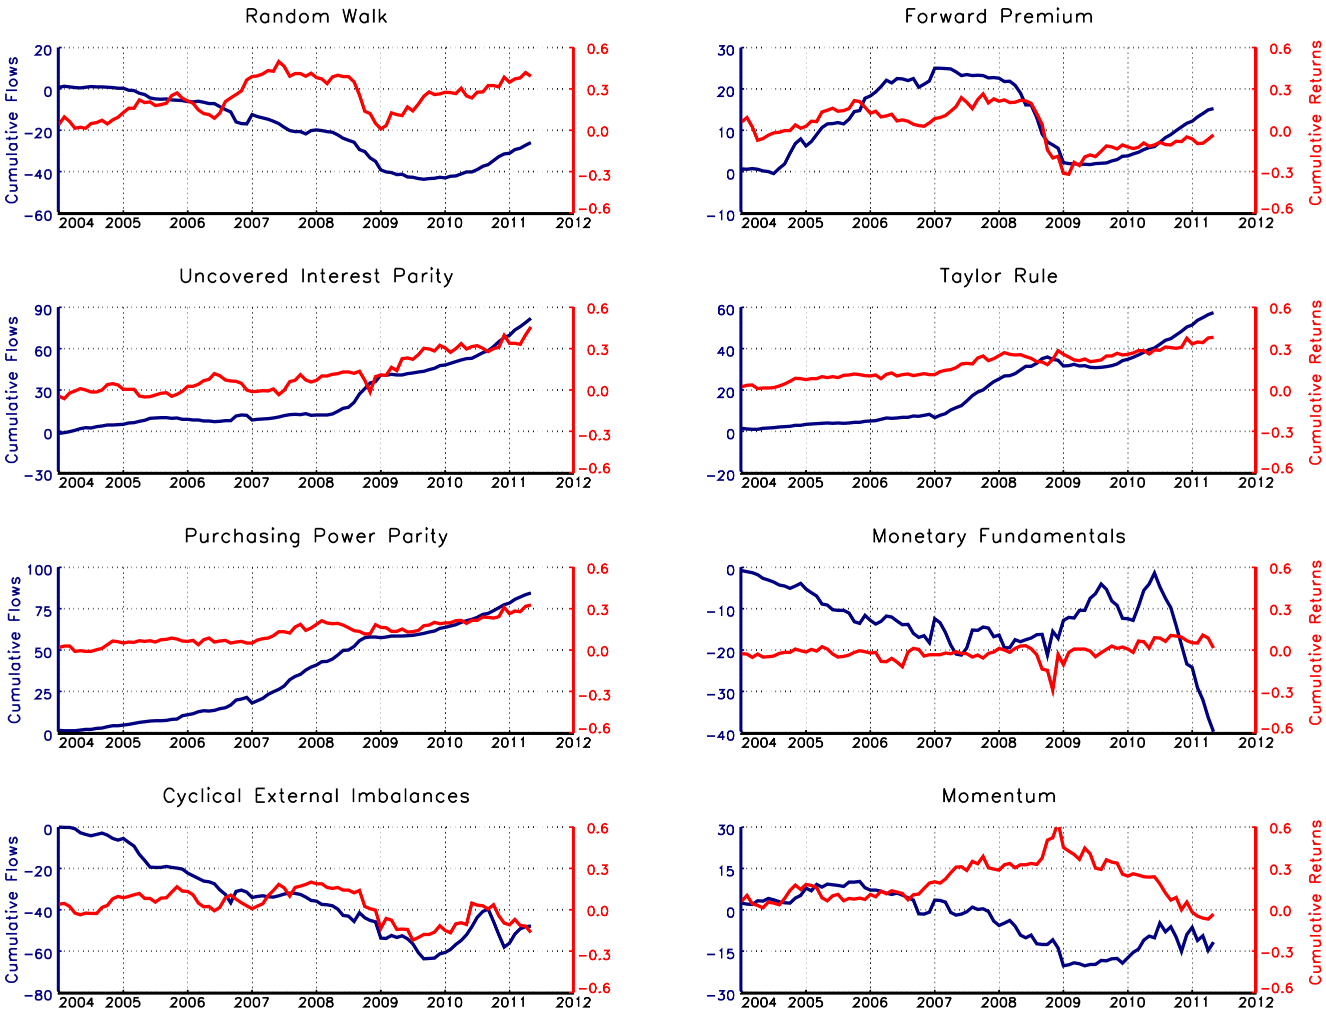

Figure 11. Cumulative Monthly Order Flow for Macroeconomic Strategies: Corporates

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 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.

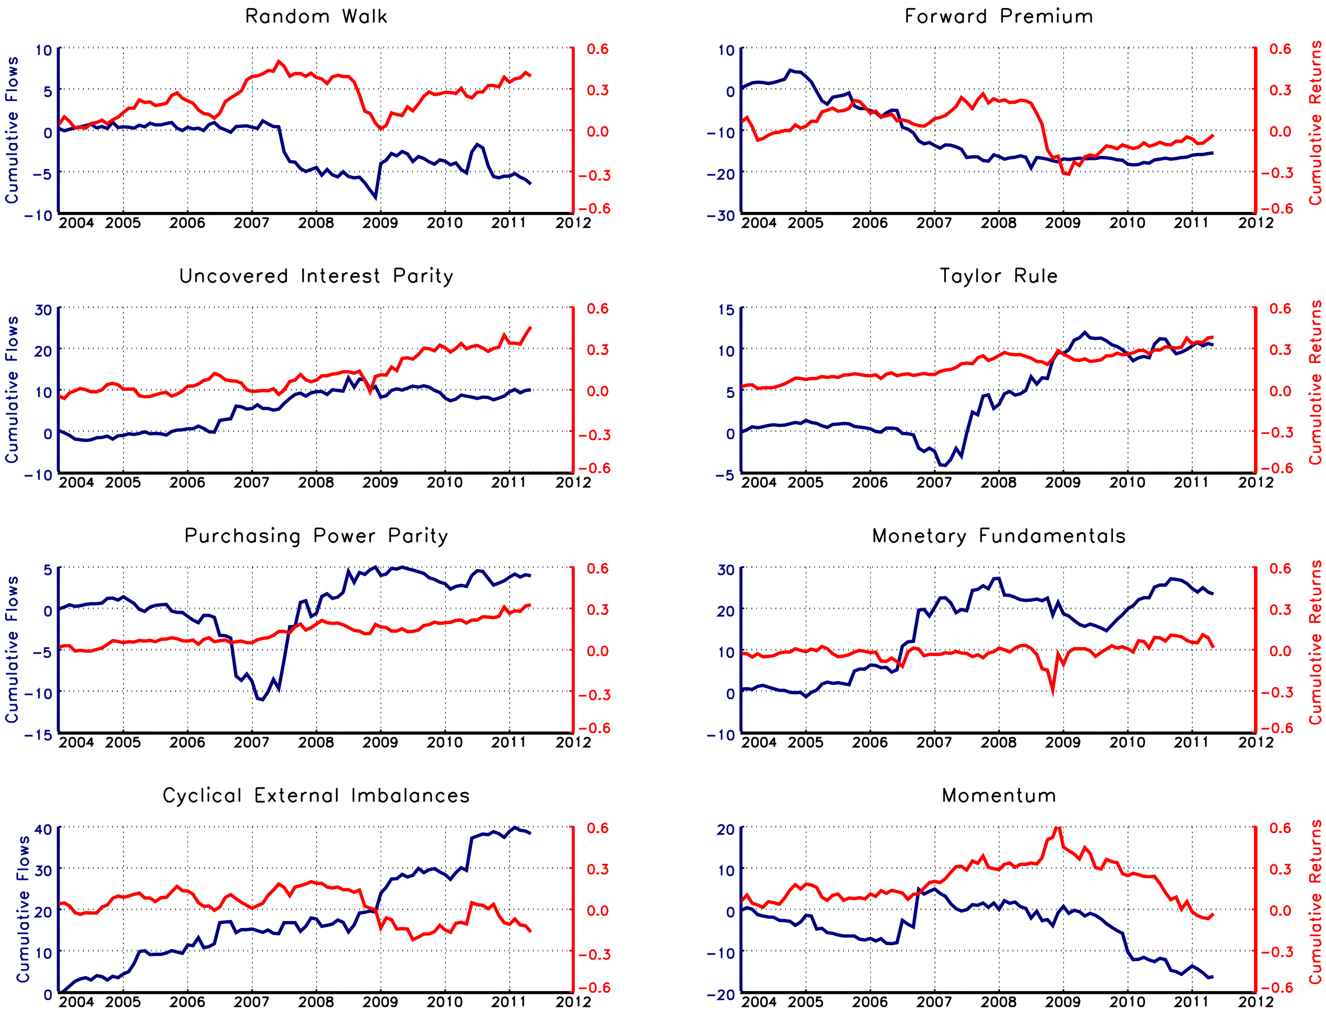

Figure 12. Cumulative Monthly Order Flow for Macroeconomic Strategies: Private Clients

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 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, F., D. Rime, and L. Sarno (2008). "Arbitrage in the Foreign Exchange Market: Turning on the Microscope," *Journal of International Economics* **76**, 237–253.

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.

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.

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

Bjonnes, G., and D. Rime (2005). "Dealer Behavior and Trading Systems in the Foreign Exchange Markets," *Journal of Financial Economics* **75**, 571–605.

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.

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.

Cerrato, M., N. Sarantis, and A. Saunders (2011). "An Investigation of Customer Order Flow in the Foreign Exchange Market," *Journal of Banking and Finance* **35**, 1892–1906.

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,

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.

DeMiguel, V., Y. Plyakha, R. Uppal, and G. Vilkov (2012). "Improving Portfolio Selection Using Option-Implied Volatility and Skewness," *Journal of Financial and Quantitative Analysis* (forthcoming).

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. (2011). *Exchange-Rate Dynamics*, Princeton: Princeton University Press.

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 (2006). “Understanding Order Flow” *International Journal of Finance and Economics* **11**, 2–23.

Evans, M.D.D., and R.K. Lyons (2007). “Exchange Rate Fundamentals and Order Flow,” Unpublished working paper, Georgetown University.

Evans, M.D.D., and R.K. Lyons (2008). “How Is Macro News Transmitted to Exchange Rates?” *Journal of Financial Economics* **88**, 26–50.

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.

Goyal, A., and A. Saretto (2009). “Cross-Section of Option Returns and Volatility,” *Journal of Financial Economics* **94**, 310–326.

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.

Kyle, A.S. (1985). “Continuous Auctions and Insider Trading,” *Econometrica* **53**, 1315–1335.

Ledoit, O., and M. Wolf (2004a). “Honey, I Shrunk the Sample Covariance Matrix,” *Journal of Portfolio Management*, **30**, 110–119.

Ledoit, O., and M. Wolf (2004b). “A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices,” *Journal of Multivariate Analysis*, **88**, 365–411.

Lobo, M.S., M. Fazel, and S. Boyd (2007). “Portfolio Optimization with Linear and Fixed Transaction Costs,” *Annals of Operations Research* **152**, 341–365.

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 (2012). "Carry Trades and Global Foreign Exchange Volatility," *Journal of Finance* **67**, 681–718.

Molodtsova, T., and D.H. Papell (2009). "Out-of-Sample Exchange Rate Predictability with Taylor Rule Fundamentals," *Journal of International Economics* **77**, 167–180.

Newey, W.K., and K.D. West (1987). "A Simple, Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* **55**, 703–708.

Obstfeld, M., and K. Rogoff (2001). "The Six Major Puzzles in International Macroeconomics: Is there a Common Cause?" *NBER Macroeconomics Annual 2000*, 339–390.

Pojarliev, M., and R.M. Levich (2008). "Do Professional Currency Managers Beat the Benchmark?" *Financial Analysts Journal* **64**, 18–32.

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.

Rogoff, K.S. (1996). "The Purchasing Power Parity Puzzle," *Journal of Economic Literature* **34**, 647–668.

Sager, M.J., and M.P. Taylor (2008). "Commercially Available Order Flow Data and Exchange Rate Movements: Caveat Emptor," *Journal of Money, Credit and Banking* **40**, 583–625.

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.