Journal of International Economics 80 (2010) 72–88

Contents lists available at ScienceDirect

Journal of International Economics

journal homepage: www.elsevier.com/locate/jie

Exchange rate forecasting, order flow and macroeconomic information

Dagfinn Rime [a,b], Lucio Sarno [c,d,*], Elvira Sojli [e]

[a] Norges Bank, Research Department, Norges Bank, P.O. Box 1179 Sentrum, 0107 Oslo, Norway
[b] Norwegian University of Science and Technology, Norway
[c] Faculty of Finance, Cass Business School, City University, 106 Bunhill Row, London EC1Y 8TZ, UK
[d] Centre for Economic Policy Research (CEPR), UK
[e] Rotterdam School of Management, Erasmus University, Department of Finance, Room T9-31, Burgemeester Oudlaan 50, P.O. Box 1738, 3000 DR, Rotterdam, The Netherlands

ARTICLE INFO

Article history:
Received 1 June 2008
Received in revised form 12 March 2009
Accepted 18 March 2009

Keywords:
Exchange rates
Microstructure
Order flow
Forecasting
Macroeconomic news

JEL classification:
F31
F41
G10

ABSTRACT

This paper adds to the research efforts that aim to bridge the divide between macro and micro approaches to exchange rate economics by examining the linkages between exchange rate movements, order flow and expectations of macroeconomic variables. The basic hypothesis tested is that if order flow reflects heterogeneous expectations about macroeconomic fundamentals, and currency markets learn about the state of the economy gradually, then order flow can have both explanatory and forecasting power for exchange rates. Using one year of high frequency data collected via a live feed from Reuters for three major exchange rates, we find that: i) order flow is intimately related to a broad set of current and expected macroeconomic fundamentals; ii) more importantly, order flow is a powerful predictor of daily movements in exchange rates in an out-of-sample exercise, on the basis of economic value criteria such as Sharpe ratios and performance fees implied by utility calculations.

1. Introduction

Following decades of failure to empirically explain and forecast fluctuations in exchange rates using traditional exchange rate determination models (Meese and Rogoff, 1983; Cheung et al., 2005; Engel et al., 2008), the recent microstructure literature has provided promising evidence, pioneered by a series of papers by Evans and Lyons (2002a, 2005a). These papers have theoretically motivated and empirically demonstrated the existence of a close link between daily exchange rate movements and order flow. Order flow is defined as the net of buyer- and seller-initiated currency transactions, and may be thought of as a measure of net buying pressure (Lyons, 2001).

In a macro–micro dichotomy of exchange rate determination, one may view the standard macro approach as based on the assumption that only common knowledge macroeconomic information matters, and the micro approach as based on the view that heterogeneous beliefs are essential to determine prices. However, given the lack of a widely accepted model for nominal exchange rates, neither of these extreme perspectives is likely to be correct. A hybrid view, as presented in the microstructure approach to exchange rates (e.g. Evans and Lyons, 2002a, 2007; Bacchetta and van Wincoop, 2006), seems more plausible. In this framework, macroeconomic information impacts on exchange rates not only directly, as in a standard macro model, but also indirectly via order flow. Order flow becomes a transmission mechanism that facilitates aggregation of dispersed price-relevant information such as heterogeneous interpretations of news, changes in expectations, and shocks to hedging and liquidity demands.

Evans and Lyons (2002a) provide evidence that order flow is a significant determinant of two major bilateral exchange rates, obtaining coefficients of determination substantially larger than the ones usually found using standard macroeconomic models of nominal exchange rates. Their results are found to be fairly robust by subsequent literature (e.g. Payne, 2003; Marsh and O'Rourke, 2005; Killeen et al., 2006). Moreover, Evans and Lyons (2005a, 2006) argue that gradual learning in the foreign exchange (FX) market can generate not only explanatory, but also forecasting power in order flow.

The finding that order flow has more explanatory power than macro variables for exchange rate behavior gives some support to the importance of heterogeneous expectations (Bacchetta and van Wincoop, 2006). However, it does not necessarily imply that order flow is the underlying determinant of exchange rates. It may well be that macroeconomic fundamentals are an important driving force for

* Corresponding author. Cass Business School, Faculty of Finance, Cass Business School, City University, 106 Bunhill Row, London EC1Y 8TZ, UK. Tel.: +44 207 0408772; fax: +44 207 0408881.
E-mail addresses: dagfinn.rime@norges-bank.no (D. Rime), lucio.sarno@city.ac.uk (L. Sarno), esojli@rsm.nl (E. Sojli).

0022-1996/$ – see front matter
doi:10.1016/j.jinteco.2009.03.005

exchange rates, but that conventional measures of expected future fundamentals are so imprecise that an order-flow "proxy" performs better in estimation. Unlike expectations measured by survey data, order flow represents a willingness to back one's beliefs with real money (Lyons, 2001).

Building on the recent success of the microstructure approach to exchange rates, a number of important hurdles remain on the route towards understanding exchange rate behavior. First, if one were willing to accept the existence of a link between order flow and exchange rate movements, economists are still awaiting for conclusive empirical evidence explaining where the information in order flow stems from. This issue is important in attempting to bridge the divide between micro and macro approaches to exchange rate economics.

Second, while the emphasis of the microstructure literature has primarily been on explaining exchange rate movements with order flow, there are only few empirical results on its forecasting power. The Meese–Rogoff finding that no available information is useful in forecasting exchange rates out-of-sample better than a naïve random walk model is robust and remains the conventional wisdom. This stylized fact implies that knowledge of the state of the economy available at a point in time is largely useless information for predicting currency fluctuations. However, if order flow does indeed reflects heterogeneous beliefs about the current and future state of the economy, and if currency markets do not discover order flow in real time but only through a gradual learning process (due to, for example, the partially decentralized nature of the FX market and its relatively low degree of transparency), then order flow should also provide forecasting power for exchange rate returns, as discussed in greater detail in the next section.

In this paper, we make progress on both these issues. We start from noting that theoretically order flow can aggregate macroeconomic information through two channels: (i) differential interpretation of news (currently available information); and (ii) heterogeneous expectations about future fundamentals. We provide evidence that the information impounded in order flow is intimately related to a broad set of macroeconomic variables of the kind suggested by exchange rate theories, as well as to expectations and changes in expectations about these fundamentals, implying that both channels suggested by theory are at work. Then, given the intermediary role of order flow for the relation between exchange rates and macroeconomic fundamentals, we investigate empirically the ability of simple microstructure models based on order flow to outperform a naïve random walk benchmark in out-of-sample forecasting.

The forecasting analysis relies on the use of economic criteria. Statistical evidence of exchange rate predictability in itself does not guarantee that an investor can earn profits from an asset allocation strategy that exploits this predictability. In practice, ranking models is useful to an investor only if it leads to tangible economic gains. Therefore, we assess the economic value of exchange rate predictability by evaluating the impact of predictable changes in the conditional FX returns on the performance of dynamic asset allocation strategies. Building on previous research by West et al. (1993), Fleming et al. (2001) and Della Corte et al. (forthcoming), we employ mean-variance analysis as a standard measure of portfolio performance and apply quadratic utility to examine whether there are any economic gains for an investor who uses exchange rate forecasts from an order flow model relative to an investor who uses forecasts from alternative specifications, including a naïve random walk model. Economic gains are evaluated mainly using two measures: the Sharpe ratio and the performance fee. The Sharpe ratio is the most common measure of performance evaluation employed in financial markets to assess the success or failure of active asset managers; it is calculated as the ratio of the average realized portfolio excess returns to their variability. The performance fee measures how much a risk-averse investor is willing to pay for switching from a portfolio strategy based on the random walk model to one which conditions on order flow. In addition, we calculate the break-even transaction cost, that is the transaction cost that would remove any economic gain from a dynamic asset allocation strategy relative to a simple random walk strategy.[1]

Using one year of data for three major exchange rates obtained from Reuters on special order, we find evidence that order flow is a powerful predictor of movements in daily exchange rates in an out-of-sample exercise, where an investor carries out allocation decisions based on order flow information. The Sharpe ratio of the order flow model is around unity and substantially higher than the Sharpe ratios delivered by alternative models, including the random walk. Furthermore, we find that a risk-averse investor would be prepared to pay high performance fees to switch from the random walk model to a model based on order flow. Consistent with leading microstructure theories, our interpretation is that order flow is a key vehicle via which fundamental information impacts on current and future prices.

The remainder of the paper is organized as follows. In the next section, we provide a brief literature review. Section 3 describes the data set and presents preliminary results on the link between order flow and exchange rates. The relation between order flow and macroeconomic fundamentals is examined in Section 4. The forecasting setup and the investor's asset allocation problem are described in Section 5, and the results on the economic value of forecasting models that condition on order flow are reported in Section 6. Section 7 concludes.

2. Related literature and motivation

The failure of conventional structural models to explain and forecast exchange rates has recently given rise to two different strands of research: one focusing on the implications of the standard present-value approach to asset pricing and the other based on the microstructure approach to the FX market. On the one hand, Engel and West (2005) demonstrate that the lack of forecastability of exchange rates using fundamentals can be reconciled with exchange rate determination theories within a rational expectations model, where the exchange rate equals the discounted present value of expected economic fundamentals. Their result is based on two assumptions: fundamentals are nonstationary processes; and the discount factor for expected fundamentals in the exchange rate equation is near unity. Under these conditions, empirical exchange rate models based on current and past macroeconomic information cannot forecast exchange rate returns, even if the model is correct, because FX returns will behave as near white noise processes. Nonetheless, Engel and West's theoretical result does not imply that fundamentals cannot forecast exchange rate fluctuations; it simply shows that lack of forecastability is not the same as rejection of the underlying model. Indeed, Engel et al. (2007) and Molodtsova and Papell (2009) find evidence that fundamentals can outperform a random walk at long horizons.

On the other hand, the microstructure literature has also taken significant steps towards understanding fluctuations in exchange rates. Evans and Lyons (2002a) propose a microstructure model that integrates public macroeconomic information and agents' heterogeneous information, where order flow serves as a mapping mechanism from dispersed information to prices. Empirically, they find that the R^2 increases from 1–5% for regressions of exchange rate changes on interest rate differentials (a proxy for public macroeconomic information) to 40–60% in regressions that use order flow to explain daily changes in exchange rates.[2] Evans and Lyons (2002b) also find that the exchange rate between two currencies is explained not only by the relevant order flow for that currency pair, but also by other currencies'

[1] There are many different ways of measuring economic gains (e.g. Leitch and Tanner, 1991), and the metrics used here are just three of them. See also Elliott and Ito (1999) and Abhyankar et al. (2005).

[2] Related papers confirming and extending these results include Payne (2003), Bjønnes et al. (2005), Killeen et al. (2006), and Berger et al. (2008).

order flows, consistent with an order-flow portfolio rebalancing approach.

At a theoretical level, Evans and Lyons (2007) formalize the notion that order flow conveys fundamental information about exchange rates in a dynamic general equilibrium model where information is first manifested at the micro (agent) level and is not symmetrically observed among agents. The model essentially combines a number of classical ingredients of the new open-economy macroeconomics literature with the insights of the FX microstructure literature, predicting an exchange rate behavior that matches several empirical facts. In a related theoretical paper, Bacchetta and van Wincoop (2006) show the existence of a close relation between order flow and exchange rates in a stylized dynamic rational expectations model. The information about macroeconomic fundamentals is assumed to be dispersed across agents, and this heterogeneity generates a larger impact of non-fundamental (e.g. hedging) trades on the exchange rate; thus a disconnect between exchange rates and fundamentals arises in the short run. However, the relation between order flow and exchange rates is strong at both short and long horizons. In essence, these papers provide significant steps towards understanding the theoretical linkages between macroeconomic fundamentals, order flow and exchange rate fluctuations in a general equilibrium setting.

Order flow may be seen as a vehicle for aggregating both differences in interpretation of news and changes in heterogeneous expectations about the future state of the economy. Starting from conventional exchange rate theories, the exchange rate can be written as the discounted present value of current and expected fundamentals:

$$s_t = (1-b)\sum_{q=0}^{\infty} b^q E_t^{\mathrm{m}} f_{t+q}, \tag{1}$$

where s_t is the log nominal exchange rate (defined as the domestic price of the foreign currency), b is the discount factor, f_t denotes the fundamentals at time t, and $E_t^{\mathrm{m}} f_{t+q}$ is the market-makers' expectation about future (q-periods ahead) fundamentals conditional on information available at time t.[3] Iterating Eq. (1) forward and rearranging terms one obtains:

$$\Delta s_{t+1} = \frac{(1-b)}{b}\left(s_t - E_t^{\mathrm{m}} f_t\right) + \epsilon_{t+1}, \tag{2}$$

where $\epsilon_{t+1} \equiv (1-b)\sum_{q=0}^{\infty} b^q \left(E_{t+1}^{\mathrm{m}} f_{t+q+1} - E_t^{\mathrm{m}} f_{t+q+1}\right)$.[4] This implies that the future exchange rate change is a function of (i) the gap between the current exchange rate and the expected current fundamentals, and (ii) a term that captures changes in expectations about fundamentals. In this setup, there is scope for order flow to reflect agents' expectations about current fundamentals (i.e. interpretations, the first term in the equation) and changes in expectations about future fundamentals that agents base their trades on (the second term in the equation). As such, the strong explanatory power of order flow for exchange rate returns can be related to a standard macroeconomic fundamentals model.

One question that arises naturally is that, if order flow reveals information about fundamentals early, it is not immediately obvious why it is so difficult to obtain cointegrating relationships between exchange rates and the same fundamentals. While we do not investigate this question in this paper, one possibility is that the long-run relation between the exchange rate and fundamentals is subject to structural breaks. For example, the theoretical model of Bacchetta and van Wincoop (2004) incorporates the fact that practitioners in the FX market regularly change the weight they attach to different economic variables. In their framework, as the market rationally searches for an explanation of the observed exchange rate change, it may attribute it to some macroeconomic indicator, which in turn becomes the "scapegoat" and influences trading strategies. The model is capable of rationalizing parameter instability in empirical exchange rate models of the kind often documented in the literature (e.g. Rossi, 2006). However, cumulative order flow, by reflecting the traders' perception of the scapegoat over time, can be related to the exchange rate in a stable cointegrating relation even though the relation between the exchange rate and fundamentals displays structural instability. A related rationalization of the lack of cointegration between exchange rates and fundamentals is offered by Chinn and Moore (2008), using an exchange rate model which is a hybrid of the conventional monetary fundamentals and the microstructure approach. In this model, shocks to preferences render the demand for money unstable and are revealed through order flow. In essence, this theory suggests a cointegrating relation involving the spot exchange rate, conventional economic fundamentals, and cumulative order flow. In this case, omitting cumulative order flow from the long-run cointegrating regression would lead to lack of cointegration due to a misspecified model.

Previous studies have found that order flow is linked to news (Dominguez and Panthaki, 2006; Evans and Lyons, 2005b, 2008; Love and Payne, 2008; Berger et al., 2008), even though the explanatory power is either not reported or documented to be very low. The role of order flow in aggregating expectations about future fundamentals has not yet been investigated in the literature. Moreover, if order flow is a proxy for the two terms in an exchange rate model of the form in Eq. (2), and the market does not discover aggregate order flow immediately, then order flow may provide forecasting power.

Given the strong contemporaneous link between exchange rates and order flow, serial correlation in order flow time series would generate forecasting power. This contemporaneous link and the serial correlation in order flow are both detected in our data set, as shown in Section 3.2 below.[5] However, two questions arise: (*i*) Why is order flow serially correlated? (*ii*) Why do dealers respond to order flow rather than simply its unexpected component? With respect to the first question, Evans and Lyons (2005a, 2006) argue that customer order flow is discovered slowly by the market, but the same argument can be applied for several other banks that have an informative clientele. Thus, it can be argued that the whole market will take at least one day to uncover the heterogeneous information embedded in order flow.[6] Moreover, order flow may be serially correlated due to order splitting across days designed not to reveal information rapidly and reduce market impact, along the lines of "stealth trading" (Chakravarty, 2001).

With respect to the second question above, it is important to note that in the canonical microstructure models of Kyle (1985) and Glosten and Milgrom (1985) the market maker only responds to unexpected order flow. However, not all assumptions of these models are true in a partially decentralized market with a complex information structure like the FX interbank market. For example, Easley and O'Hara (1992) examine the adverse selection problem that arises from the repeated trades of informed traders. Their model shows that trading volume affects the speed of price adjustment to information and that the efficiency of price adjustment to new information depends on the specific market structure. In market structures that have richer

[3] The model is adapted from Engel and West (2005) who use market expectations about macroeconomic fundamentals, not the expectations of market makers.

[4] Usually present-value models of this kind assume that $Ef_t = f_t$, i.e. that current fundamentals are observable without error in real time. However, in practice, macroeconomic data are not available in real time, since most macro data reported at time t relate to values for a previous month or quarter and tend to contain measurement errors (see Faust et al., 2003; Sarno and Valente, 2009).

[5] Breedon and Vitale (2005) also report serial correlation in Reuters order flow over a different sample period.

[6] Note that even custodian banks, which record order flows for a large proportion of the FX market, typically release data with significant lags in order to protect clients' confidentiality and meet compliance requirements. For example, State Street, a major custodian bank, releases FX order flow data with a 4-day delay, implying that the learning process discussed above may take several days.

informational problems than described in the canonical microstructural models, asset prices do not necessarily reflect all new information instantaneously.

In light of the above considerations, several studies have investigated the existence of forecasting power in order flow for exchange rate returns. However, the evidence is scant and mixed. On the one hand, Evans and Lyons (2005a, 2006) use six years of proprietary disaggregated customer data on US dollar–euro from Citigroup and find that the forecasts based on an order-flow model outperform the random walk at various forecast horizons ranging from 1 to 20 trading days. On the other hand, Danielsson et al. (2002) and Sager and Taylor (2008) find no evidence of better forecasting ability in order flow models relative to a random walk benchmark for several major exchange rates and different forecast horizons. Similarly, Killeen et al. (2006) estimates a partial adjustment (error correction) model for the effect of order flow on exchange rate returns and finds that the speed of convergence to the long-run equilibrium is very fast, implying that the predictive information content in order flow decays rapidly. Hence, the forecasting results obtained by Evans and Lyons (2005a, 2006) are waiting to be tested by other studies and with alternative data sources, especially because their data is not available, given their confidential nature. It is also important to note that all these studies carry out forecast accuracy tests using conventional statistical methods. As mentioned earlier, we move away from standard statistical metrics of forecast evaluation and rely on measures of the economic value of order flow in an asset allocation setting.

3. Data and preliminaries

3.1. Data sources

The FX market is by far the largest financial market, with a daily turnover of US dollar (USD) 3210 billion, a third of which is in spot transactions. Electronic brokers have become the preferred means of settling trades, and 50–70% of turnover in the major currency pairs is settled through the two main electronic platforms, Reuters and Electronic Brokerage System (EBS) (Galati and Melvin, 2004).[7] Most previous studies in exchange rate microstructure have used data from the early phase of electronic brokers in this market (before 2000), with the exception of Berger et al. (2008).

This paper uses interdealer data for three major exchange rates: USD *vis-à-vis* the euro, the UK sterling and the Japanese yen (hereafter EUR, GBP and JPY respectively), for the sample period from February 13, 2004 to February 14, 2005. The data set includes all best ask and bid quotes as well as all trades in spot exchange rates. The data is obtained from Reuters trading system (D2000-2) on special order and collected via a continuous live feed.[8] The Bank for International Settlements (2005) estimates that trades in these currencies constituted up to 60% of total FX transactions in 2004, the period we are investigating; hence, trading in the three currency pairs studied here comprises a substantial part of the FX market. However, it is important to note that, while Reuters is the platform where most of the GBP trades take place, EBS has the highest share of trades in EUR and JPY. The smaller coverage of Reuters for EUR and JPY relative to EBS has ambiguous implications for the empirical analysis of the linkage between order flow and exchange rates. On the one hand, one would expect that the platform with the largest coverage of an exchange rate provides a more precise signal of trading in the overall FX market and, therefore, a more powerful measure of order flow in terms of information content. On the other hand, microstructure theory provides counter-arguments to this view. Notably, Admati and Pfleiderer (1988) show that uninformed traders would bunch together on the most liquid platform to reduce their transaction costs—this is termed “clumping” behavior (O'Hara, 1995). This suggests that it is not obvious that the information content of order flow in the most liquid or larger platform (EBS for EUR and JPY, Reuters for GBP) is superior to the less liquid or smaller platform (EBS for GBP, Reuters for EUR and JPY).[9]

Daily data are constructed from tick data and include the most active part of the trading day between 7:00 and 17:00 GMT. In addition, weekends, holidays and days with unusually low or no trading activity (due to feed failures) are excluded. Using daily data allows to filter out transitory liquidity effects and to focus on a horizon that is relevant to market participants. Order flow, Δx_t, is measured as the aggregated difference between the number of buyer-initiated and seller-initiated transactions for the foreign (base) currency from 7:00 to 17:00 GMT; positive (negative) order flow implies net foreign currency purchases (sales).[10] The daily exchange rate is expressed as the USD value of one unit of foreign currency; the daily exchange rate return, Δs_t, is calculated as the difference between the log midpoint exchange rate at 7:00 and 17:00 GMT, whereas in the forecasting exercise it is defined as the difference between the midpoint rate at 17:00 of day t and 17:00 of day $t-1$. The former definition matches exactly the definition of order flow and is useful for contemporaneous regressions, whereas the latter is more appropriate for the forecasting exercise where the investor is assumed to forecast exchange rates one-day-ahead on the basis of information that is available at 17:00 GMT on day t to forecast exchange rates at 17:00 on day $t+1$.

It is important to note that the data used here are different from the customer order flow data employed in some of the papers cited earlier (e.g. Evans and Lyons, 2005a). While the customer order flow data are proprietary (hence not publicly available), the tick-by-tick Reuters data can be observed directly on a Reuters dealing screen. However, it is hardly possible to define this data as public since Reuters does not generally provide historical data on order flow. In essence, utilization of this data first requires a special order and authorization to download via a live feed, then careful analysis is necessary to aggregate the data from tick frequency to generate signed daily order flow data. This cumbersome process demands both (expensive) special authorizations and IT resources that constitute a serious barrier to data gathering for the uninitiated in this area of research. In this sense, we would argue that Reuters order flow data do not constitute public information in the sense that they are not simply available by any data provider in real time.

The interest rates used are the overnight LIBOR fixings for euro, UK sterling, US dollar and the spot/next LIBOR fixing for Japanese yen, obtained from *EcoWin*. Data on economic fundamentals is provided from the Money Market Survey (MMS), carried out by InformaGM. The data set includes values for expected, announced and revised macroeconomic variables. Market participants' expectations on macroeconomic fundamentals are collected weekly and aggregated on Thursday the week prior to the announcement week. Note that because information on macroeconomic fundamentals is published with a lag, their values pertain to the month or quarter prior to the current one. We have data for announcements over the period from February 13, 2004 to February 14, 2005 for EMU, the UK and the US.[11]

[7] For a detailed description of the structure of the FX market and electronic trading platforms, see Lyons (2001) and Rime (2003).

[8] Reuters generally provides only data on the number, not the volume, of trades, but this should not influence the empirical analysis and results. Bjønnes and Rime (2005) and Killeen et al. (2006) show that analysis based on trade size and number of trades is not qualitatively different.

[9] It is unfortunate that studies of FX order flow generally do not use both Reuters and EBS data, which would allows us to learn more about the implications of these issues on the empirical front. The exception to this caveat is the work of Breedon and Vitale (2005), which uses data from both EBS and Reuters for one exchange rate (EUR) over a four-month period in 2000.

[10] In a limit order book like Reuters, the initiator is the one that consumes liquidity services and pays half of the spread in order to make a transaction. Liquidity providers use limit orders; liquidity consumers use market orders.

[11] The full list of variables for which we have announcements is available upon request. Note that this data is not available for Japan over the sample period.

3.2. Preliminary analysis

Summary statistics for daily exchange rate returns and order flows are reported in Panel A of Table 1. The properties of exchange rate returns are similar across currencies: mean returns are very close to zero and standard deviations are large and of similar magnitude across currencies. The mean of daily order flows is positive, implying positive demand for foreign currencies in the sample period under investigation. Standard deviations are fairly large, allowing for negative order flows and positive demand for the USD in certain periods of time during the sample.

Panel B of Table 1 shows that there is high positive correlation among exchange rate returns, partly due to the common denomination against the USD. The highest correlation is observed between EUR and GBP. Correlations between currency pairs and the relevant order flows are high, above 0.4, and those with other currency pairs' order flows are also sizable.

Panel C of Table 1 exhibits the first-order serial correlation of the order flow time series (ranging from 8 to 14%). Also there is some evidence of sizable correlations from order flow in one currency and the next-day order flow in another currency. This is the case for the order flow of JPY, which appears to be correlated with next-day order flow in both EUR and GBP. In turn, the strong correlations reported in Panels B–C suggest that system estimation of regressions involving exchange rates and order flow may be superior to single-equation models (see Evans and Lyons, 2002b).

Hence, in order to allow for cross-currency effects of order flow, we use the seemingly unrelated regressions (SUR) method to estimate:

$$\boldsymbol{P}_t = \boldsymbol{C} + \boldsymbol{B}\boldsymbol{X}_t + \boldsymbol{V}_t, \tag{3}$$

where $\boldsymbol{P}_t$ is the 3×1 vector of exchange rate changes, $\boldsymbol{P}_t = [\Delta s_t^{\mathrm{EUR}}, \Delta s_t^{\mathrm{GBP}}, \Delta s_t^{\mathrm{JPY}}]'$; $\boldsymbol{X}_t$ is the 3×1 vector of order flows, $\boldsymbol{X}_t = [\Delta x_t^{\mathrm{EUR}}, \Delta x_t^{\mathrm{GBP}}, \Delta x_t^{\mathrm{JPY}}]'$; $\boldsymbol{B}$ is the 3×3 matrix of order flow coefficients; $\boldsymbol{C}$ is the vector of constant terms; and $\boldsymbol{V}_t$ is the vector of error terms. The results in Table 2 show that estimation of model (3) yields very strong explanatory power (R^2) for all currencies. 'Own' order flow (that is the order flow of the currency pair on the left-hand side of the equation) has a significant positive coefficient for all the exchange rate movements, but the cross-currency order flows also have significant effects on exchange rate returns, consistent with the above cited studies. The Wald test statistic strongly rejects the null hypothesis that the order flow coefficients in each regression are jointly equal to zero.

Table 1
Preliminary data analysis.

	$\Delta s_t^{\mathrm{EUR}}$	$\Delta s_t^{\mathrm{GBP}}$	$\Delta s_t^{\mathrm{JPY}}$	$\Delta x_t^{\mathrm{EUR}}$	$\Delta x_t^{\mathrm{GBP}}$	$\Delta x_t^{\mathrm{JPY}}$
Panel A. Descriptive statistics						
Mean	−0.003	−0.03	−0.02	23.18	83.00	2.21
Std. dev.	0.53	0.49	0.51	124.90	149.20	19.50
Skewness	0.29	0.002	−0.03	0.26	0.45	−0.31
Kurtosis	4.35	3.11	4.59	3.64	3.41	4.46
Panel B. Contemporaneous correlations						
$\Delta s_t^{\mathrm{EUR}}$	1.00					
$\Delta s_t^{\mathrm{GBP}}$	0.70	1.00				
$\Delta s_t^{\mathrm{JPY}}$	0.46	0.46	1.00			
$\Delta x_t^{\mathrm{EUR}}$	0.65	0.53	0.43	1.00		
$\Delta x_t^{\mathrm{GBP}}$	0.35	0.42	0.30	0.38	1.00	
$\Delta x_t^{\mathrm{JPY}}$	0.20	0.28	0.49	0.23	0.15	1.00
Panel C. Order flows correlations						
$\Delta x_{t-1}^{\mathrm{EUR}}$				0.13	0.12	0.00
$\Delta x_{t-1}^{\mathrm{GBP}}$				−0.01	0.08	−0.01
$\Delta x_{t-1}^{\mathrm{JPY}}$				0.06	0.22	0.14

Preliminary analysis for the period 2/13/2004–2/14/2005. Δs_t^j is the daily exchange rate return from 7:00 to 17:00 GMT, and Δx_t^j is the daily order flow (positive for net foreign currency purchases) accumulated between 7:00 and 17:00 GMT, for each exchange rate j: US dollar/euro (EUR), US dollar/UK sterling (GBP) and US dollar/Japanese yen (JPY). Panel A presents descriptive statistics for exchange rate returns and order flows. The means and standard deviations for exchange rate returns are expressed in percentage points. Panel B exhibits sample contemporaneous correlations among exchange rate returns and order flows. In Panel C the diagonal elements show the autocorrelations, i.e. $corr(\Delta x_t^j, \Delta x_{t-1}^j)$. The off-diagonal elements are the lagged cross-correlations, i.e. $corr(\Delta x_t^j, \Delta x_{t-1}^i)$, where $i \neq j$.

Table 2
Order flow model.

	$\Delta s_t^{\mathrm{EUR}}$		$\Delta s_t^{\mathrm{GBP}}$		$\Delta s_t^{\mathrm{JPY}}$	
$\Delta x_t^{\mathrm{EUR}}$	**2.52**	**(8.91)**	**1.59**	**(5.68)**	**1.18**	**(4.12)**
$\Delta x_t^{\mathrm{GBP}}$	**0.41**	**(1.78)**	**0.85**	**(3.74)**	**0.45**	**(1.93)**
$\Delta x_t^{\mathrm{JPY}}$	1.22	(0.72)	**4.18**	**(2.47)**	**10.10**	**(5.83)**
Wald test	[0.00]		[0.00]		[0.00]	
R^2	0.44		0.38		0.36	

SUR estimates of model (3) for the period 2/13/2004–2/14/2005. Δs_t^j is the daily exchange rate return from 7:00 to 17:00 GMT, and Δx_t^j is the daily order flow (positive for net foreign currency purchases), accumulated between 7:00 and 17:00 GMT, for each exchange rate j: EUR, GBP and JPY. The coefficients of the explanatory variables are expressed in percentage terms for a purchase of a thousand units of the base currency. t-statistics are shown in parenthesis. Coefficients in bold are significant at least at the 10% level of significance. The Wald test presents the p-value (in square brackets) for the joint null hypothesis that all order flow coefficients are equal to zero. All equations are estimated with a constant, which is not reported to conserve space.

4. Order flow and macroeconomic fundamentals

In this section, we examine the link between macroeconomic information and order flow using the standard present-value exchange rate model:

$$\Delta s_{t+1} = \frac{(1-b)}{b}\left(s_t - E_t^{\mathrm{m}} f_t\right) + \epsilon_{t+1}, \tag{4}$$

where $\epsilon_{t+1} \equiv (1-b)\sum_{q=0}^{\infty} b^q \left(E_{t+1}^{\mathrm{m}} f_{t+q+1} - E_t^{\mathrm{m}} f_{t+q+1}\right)$. As discussed previously, in this model order flow may capture current fundamentals information (the first term in Eq. (4)) and changes in expectations about future fundamentals (the second term in Eq. (4)). We investigate empirically both links between order flow, expectations and news.

4.1. The link between order flow and news

First, we investigate whether news explain order flow. News may trigger different interpretations for the equilibrium exchange rate and induce agents to trade, so that news could explain order flow fluctuations. Put another way, heterogeneous interpretations of the impact of news on the exchange rate leads market makers to make inferences about the equilibrium exchange rate from aggregate order flow.

News are calculated as $d_{n,t} = \frac{a_{n,t-k} - E_{t-l}a_{n,t-k}}{\sigma_n}$, where $a_{n,t-k}$ is the actual value of indicator n (say GDP, inflation, etc.) at time t pertaining to the indicator at time $t-k$; k is a week, month or quarter; $E_{t-1}a_{n,t-k}$ is the expected value of indicator n formed at time $t-l$ (the survey expectation), where l ranges between 2 and 6 trading days; and σ_n is the sample standard deviation for indicator n.[12] For each order flow series, we estimate the regression

$$\Delta x_t = \phi_0 + \sum_{n=1}^{N} \phi_n d_{n,t} + u_t \tag{5}$$

using OLS. In theory, positive news about a country ought to lead to an appreciation of its currency, but it is important to note that this does not necessarily mean that order flow has to be positively related to

[12] Ideally, we would like to have expectations on fundamentals just before the announcement time, since expectations can change in a week. This data, however, is not available to us.

good news. In a hybrid model with rational expectations and order flow, it is possible that the initial reaction of exchange rates to news fully captures the news or even over-reacts to it; in this case subsequent trading (order flow) may even be negatively related to positive news. In other words, the sign of the relation between news and order flow is ambiguous since it will depend on the extent to which the exchange rate adjusts directly in response to the news (see Evans and Lyons, 2008 for a discussion of this issue).

The results from estimating Eq. (5) are presented in Table 3. The estimated coefficients reported are statistically significant at least at the 10% level, suggesting that news are an important determinant of order flow. Moreover, our results suggest that demand for a currency is stronger in response to good news, i.e. positive news on the US economy are associated with a decrease in order flow (stronger demand for USD), whereas positive news on foreign economies are associated with an increase in order flow (stronger demand for the base currency), consistent with related evidence in Love and Payne (2008). Such a response may be explained by an initial under-reaction of the exchange rate to news.

The news that have the highest explanatory power for order flow are similar to those that Andersen et al. (2003) find significant in explaining exchange rate fluctuations at the intraday frequency around macroeconomic announcements. These include, for example, news related to economic activity, inflation, non-farm payroll and confidence indicators. Macroeconomic news can explain up to 15% of the daily fluctuations in order flow.[13]

The microstructure approach predicts that information impacts on exchange rates both directly and indirectly via order flow (Lyons, 2001; Evans and Lyons, 2008). We assess these two channels by regressing exchange rate returns on both news and order flow, using the macroeconomic news that explain order flow as explanatory variables. The results (not reported to conserve space), show that macroeconomic news can explain fluctuations in the daily exchange rates (direct channel), and that there is an additional role for order flow (indirect channel). We find that the addition of order flow significantly increases the explanatory power for exchange rate fluctuations, as compared to news alone. Furthermore, the combined explanatory power of order flow and news appears to be higher than that of order flow alone. Thus, there is evidence of a dual impact of macroeconomic news on exchange rates, both direct and indirect via order flow.

Table 3
Contemporaneous effect of news on order flow.

Announcement	EUR	GBP	JPY
	Estimated parameters		
US			
Chicago PMI	−115.40		
Construction spending			−4.02
Consumer confidence index			−15.30
Consumer credit	136.95		
Consumer price index		−171.41	
Durable goods orders		−134.95	
GDP advance	−114.40		
GDP preliminary	−69.59		
Housing starts		−5.05	
Initial unemployment claims			4.21
Michigan sentiment (final)	−98.90		−5.36
Nonfarm payroll employment	−90.42	−88.23	
Trade balance	−59.74		−6.51
Unemployment rate	90.69		
EMU			
Consumer confidence balance	149.03		
Consumer price index	160.52		
Industrial production (yoy)	62.78		
Labor costs	101.08		
Retail sales (mom)	82.30		
Sentiment index	179.60		
UK			
GDP provisional (qoq)		267.86	
Trade balance		93.42	
R^2	0.15	0.15	0.03
Serial correlation	[0.01]	[0.76]	[0.07]
Heteroskedasticity	[0.99]	[0.99]	[0.99]

OLS regression of order flow (1000 net purchases) for EUR, GBP and JPY on contemporaneous news: $\Delta x_t = \phi_0 + \sum_{n=1}^{N} \phi_n d_{n,t} + u_t$. News are the difference between the actual value ($a_{n,t-k}$) of the macroeconomic indicator minus its expected value ($E_{t-l}a_{n,t-k}$), standardized by the sample standard deviation (σ_n), $d_{n,t} = \frac{a_{n,t-k} - E_{t-l}a_{n,t-k}}{\sigma_n}$. The regression is estimated on all the indicators available, for the period 2/13/2004–2/14/2005. Serial correlation presents the *p*-values for the Breusch–Godfrey Lagrange multiplier tests for first-order residual serial correlation. Heteroskedasticity shows the *p*-values for the White first-order conditional heteroskedasticity test with cross terms in the residuals. Only variables significant at least at the 10% level using heteroskedasticity- and autocorrelation-consistent standard errors are reported. All equations are estimated with a constant, which is not reported in order to conserve space. The total number of observations for each of the three equations is 263, for the period 2/13/2004–2/14/2005.

4.2. The link between order flow and expectations

In this sub-section, we examine the hypothesis that order flow aggregates changes in expectations. Given that the survey expectations about fundamentals are collected and published each Thursday before the announcement week, starting from the survey expectation day (i.e. Thursday), agents can revise their expectations from $E_{t-1}a_{n,t-k}$ to $E_t a_{n,t-k}$ and trade on these expectation changes. This implies that, in principle, revisions in expectations between the day of collecting survey expectations until the day of the macroeconomic announcement may be reflected in order flow.

This hypothesis can be tested by using the sum of order flows between Thursday and the announcement day to explain news, $d_{n,t}$:

$$d_{n,t} = \theta_0 + \theta_1 sumx_t + \epsilon_t \tag{6}$$

where $sumx_t = \sum_{h=0}^{l-1} \Delta x_{t-h}$ is the sum of order flow for each currency from the day of forming the survey expectation (Thursday) to the announcement day for indicator *n*; *l* varies between 2 and 6; and ε_t is the error term. For example, if the actual industrial production figure for EMU (US) is higher than the survey expected value, then rational expectation revisions prior to the news release will lead to more demand for EUR (USD). In turn, the order flow coefficient for EUR is expected to be positive (negative). The opposite will occur for variables whose impact on the economy is considered bad news, e.g. unemployment, etc. The table below indicates how order flow and the coefficients of the above equation are expected to behave if order flow is taken to be a proxy for the change in expectations between the survey and the announcement days, for US and foreign (F) news:

	"Good" news		"Bad" news	
$a^{US}_{n,t-k} > E_{t-1}a^{US}_{n,t-k}$	$\Delta x^F < 0$	$\theta_1 < 0$	$\Delta x^F > 0$	$\theta_1 > 0$
$a^{US}_{n,t-k} < E_{t-1}a^{US}_{n,t-k}$	$\Delta x^F > 0$	$\theta_1 < 0$	$\Delta x^F < 0$	$\theta_1 > 0$
$a^{F}_{n,t-k} > E_{t-1}a^{F}_{n,t-k}$	$\Delta x^F > 0$	$\theta_1 > 0$	$\Delta x^F < 0$	$\theta_1 < 0$
$a^{F}_{n,t-k} < E_{t-1}a^{F}_{n,t-k}$	$\Delta x^F < 0$	$\theta_1 > 0$	$\Delta x^F > 0$	$\theta_1 < 0$

Given that estimation of above equation would be based on a small number of observations, regression estimates are likely to suffer from small sample bias. Since we are primarily interested in the sign of the relation between order flow and news, Figs. 1–5 show the scatter plots of news against aggregated order flow to gauge whether the relation is positive or negative. The scatter plots suggest that order flow and news exhibit the expected relation for most indicators in all the

[13] Note that inflation can be, in theory, associated with both an appreciation and a depreciation of the domestic currency. It is often found that the US dollar appreciates when higher-than-expected inflation is announced (e.g. Engel et al., 2007), although a standard monetary model would imply a depreciation.

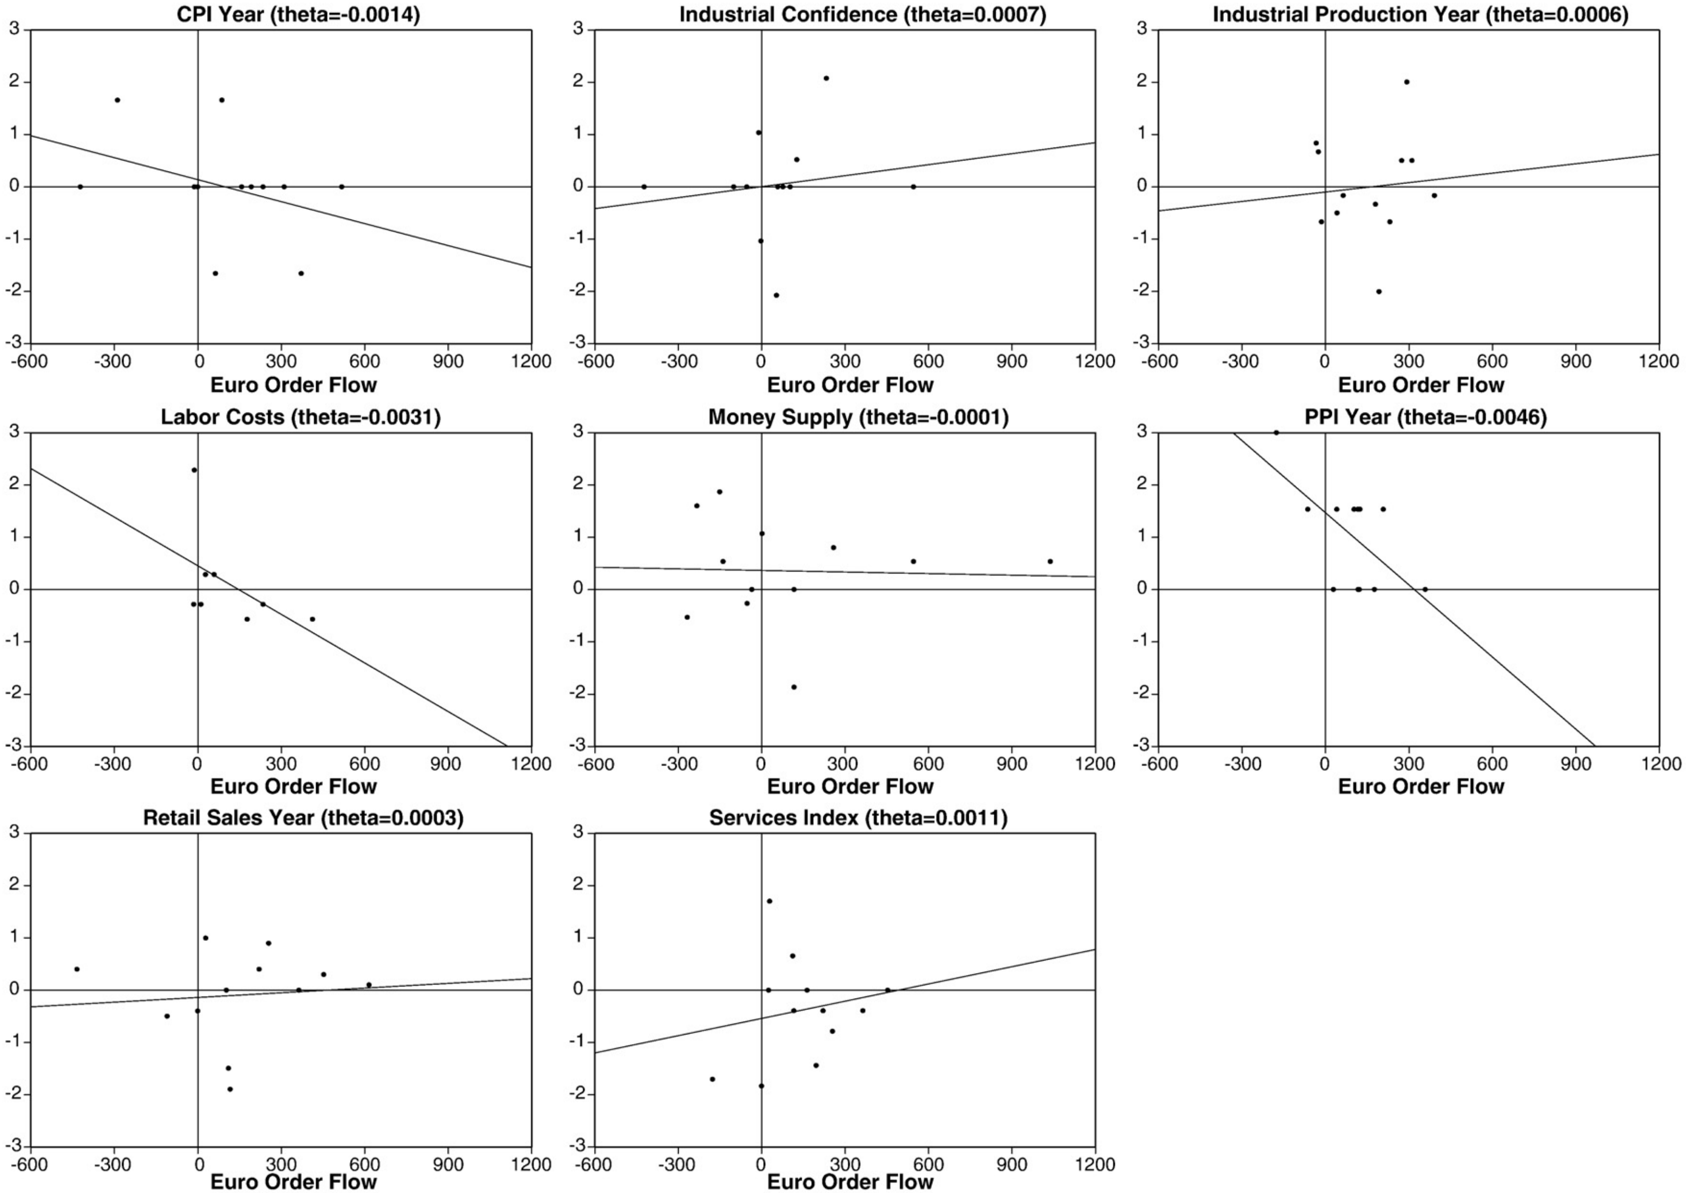

Fig. 1. EMU news and order flow. The figure presents the scatter plot of cumulated EUR order flow (*sumx*, on the horizontal-axis) and the standardized expectations gap (on the vertical-axis) for EMU news. The line describes the linear relation between order flow and news. Theta is the elasticity of the standardized expectations gap to one unit of order flow.

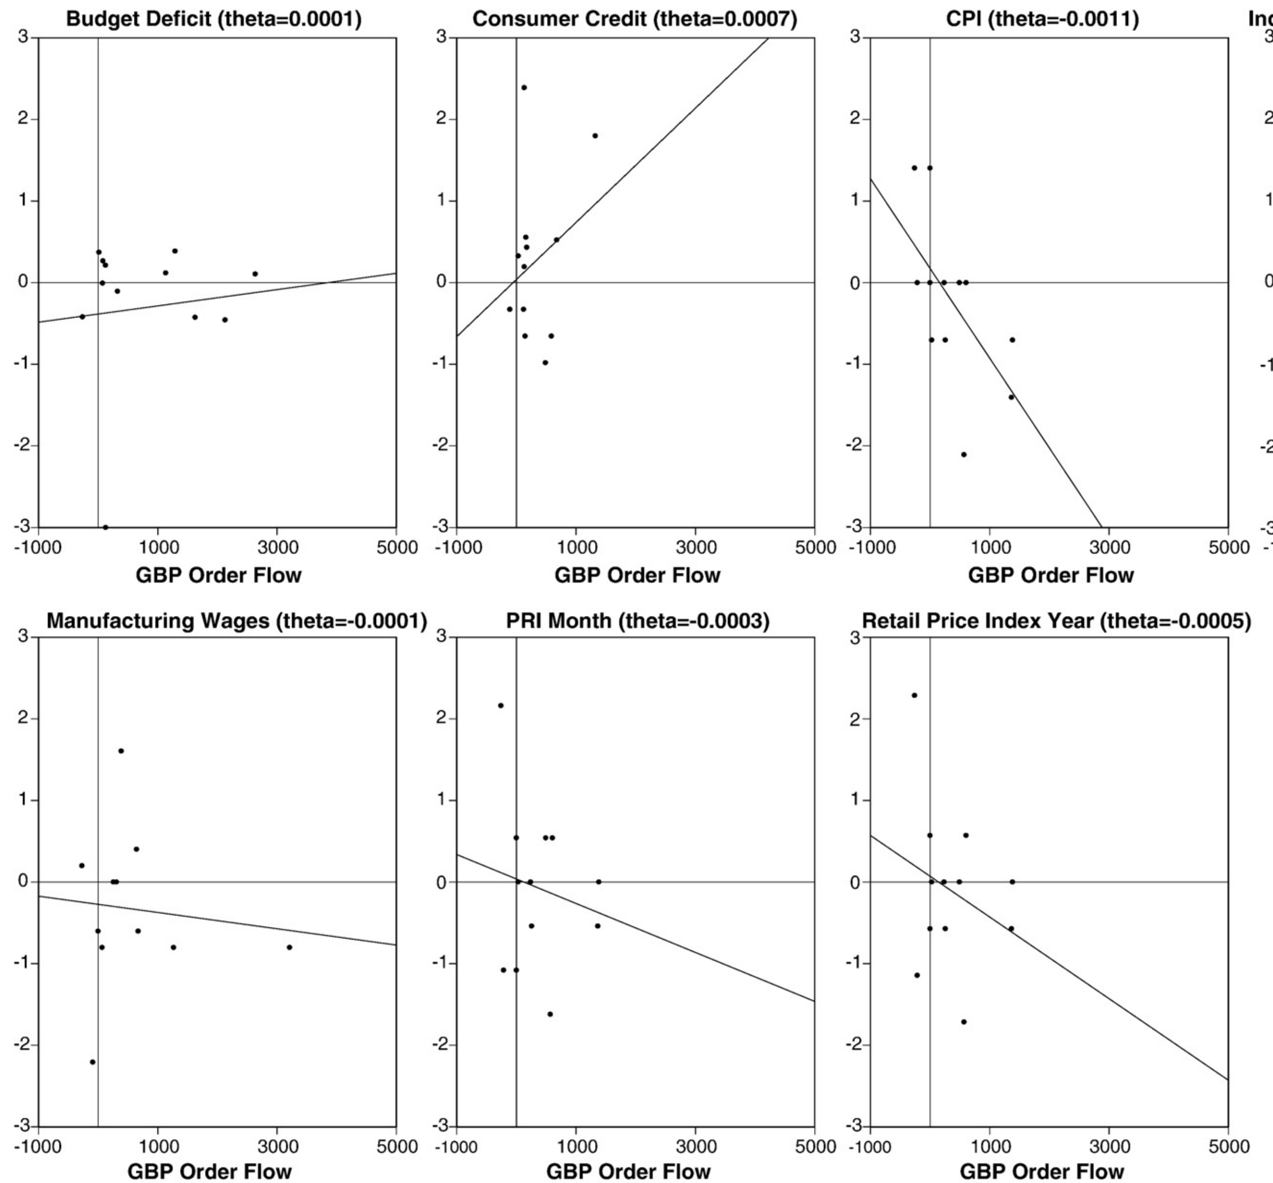

Fig. 2. UK news and order flow. The figure presents the scatter plot of cumulated GBP order flow (*sumx*, on the horizontal-axis) and the standardized expectations gap (on the vertical-axis) for UK news. The line describes the linear relation between order flow and news. Theta is the elasticity of the standardized expectations gap to one unit of order flow.

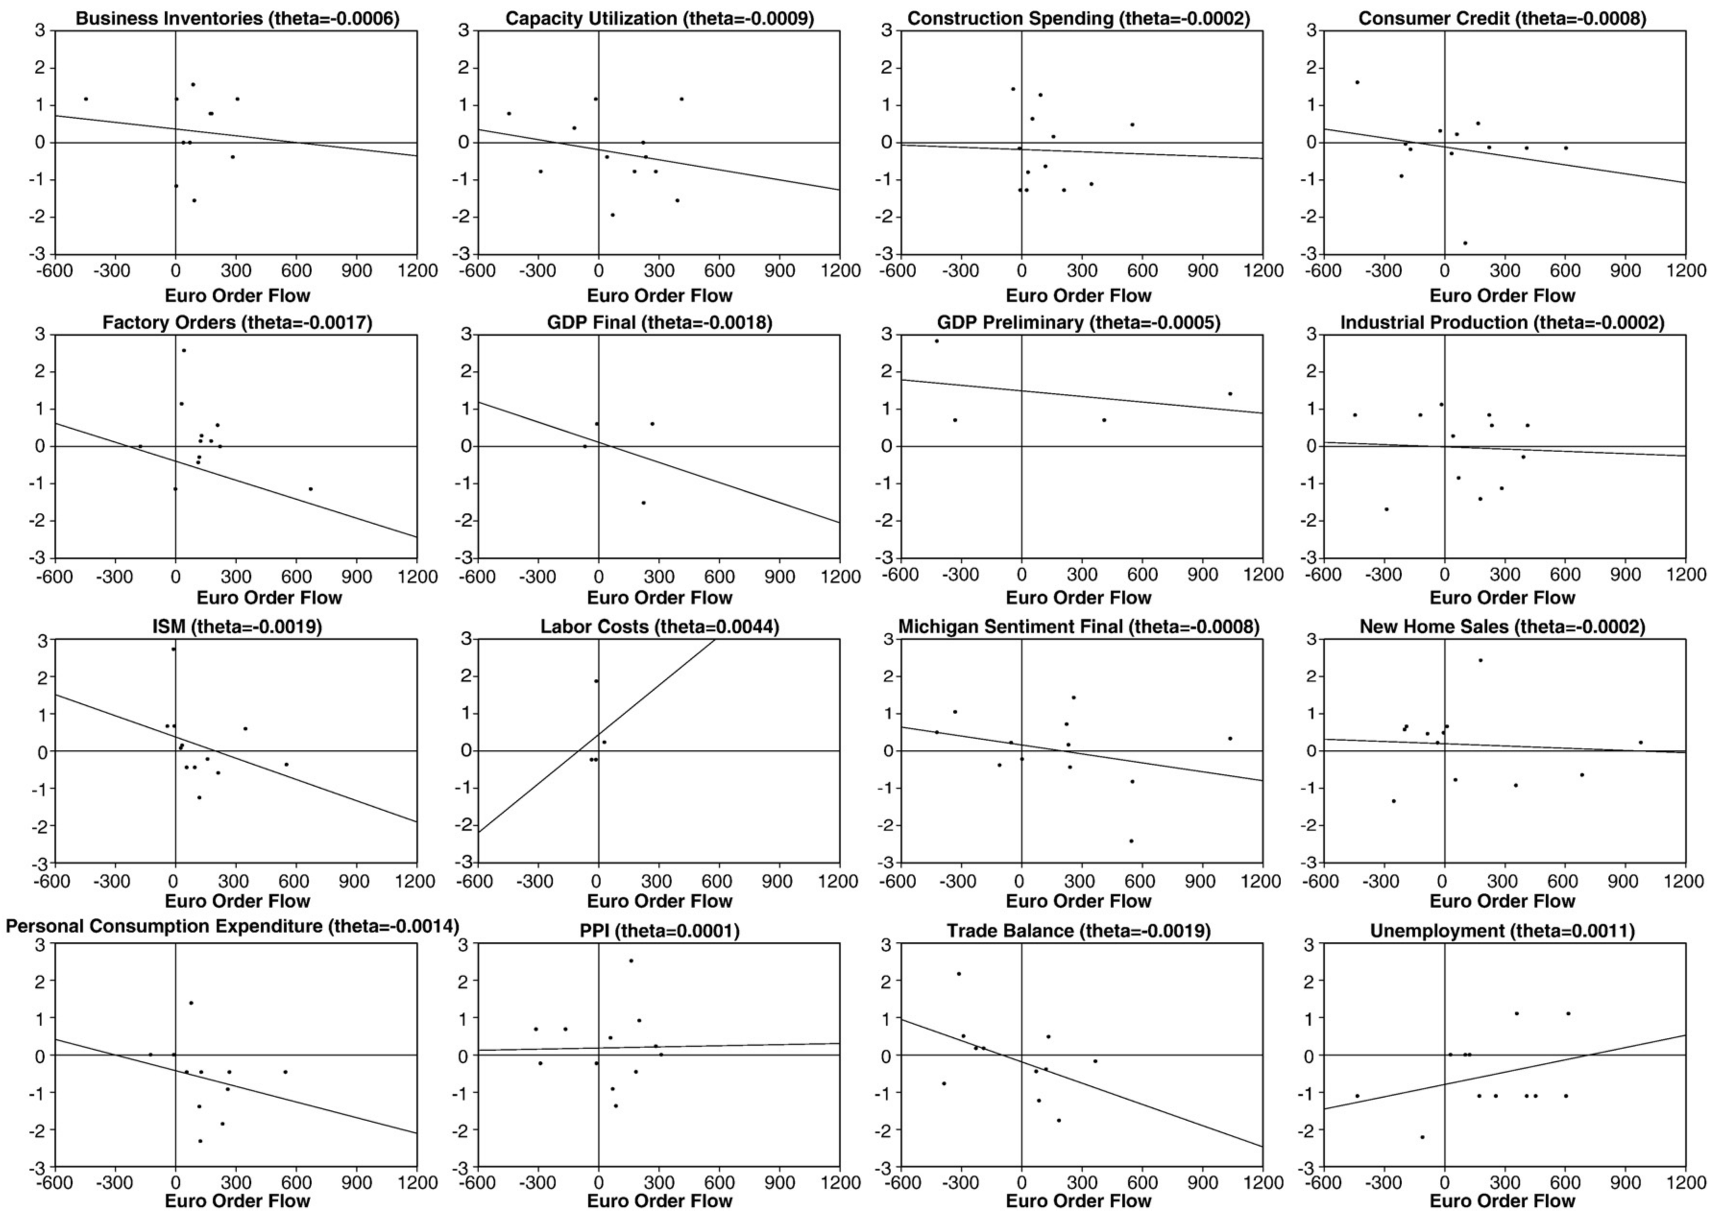

Fig. 3. US news and EUR order flow. The figure presents the scatter plot cumulated EUR order flow (*sumx*, on the horizontal-axis) and standardized expectations gap (vertical-axis) for the US news. The line describes the linear relation between order flow and news. Theta is the elasticity of the standardized expectations gap to one unit of order flow.

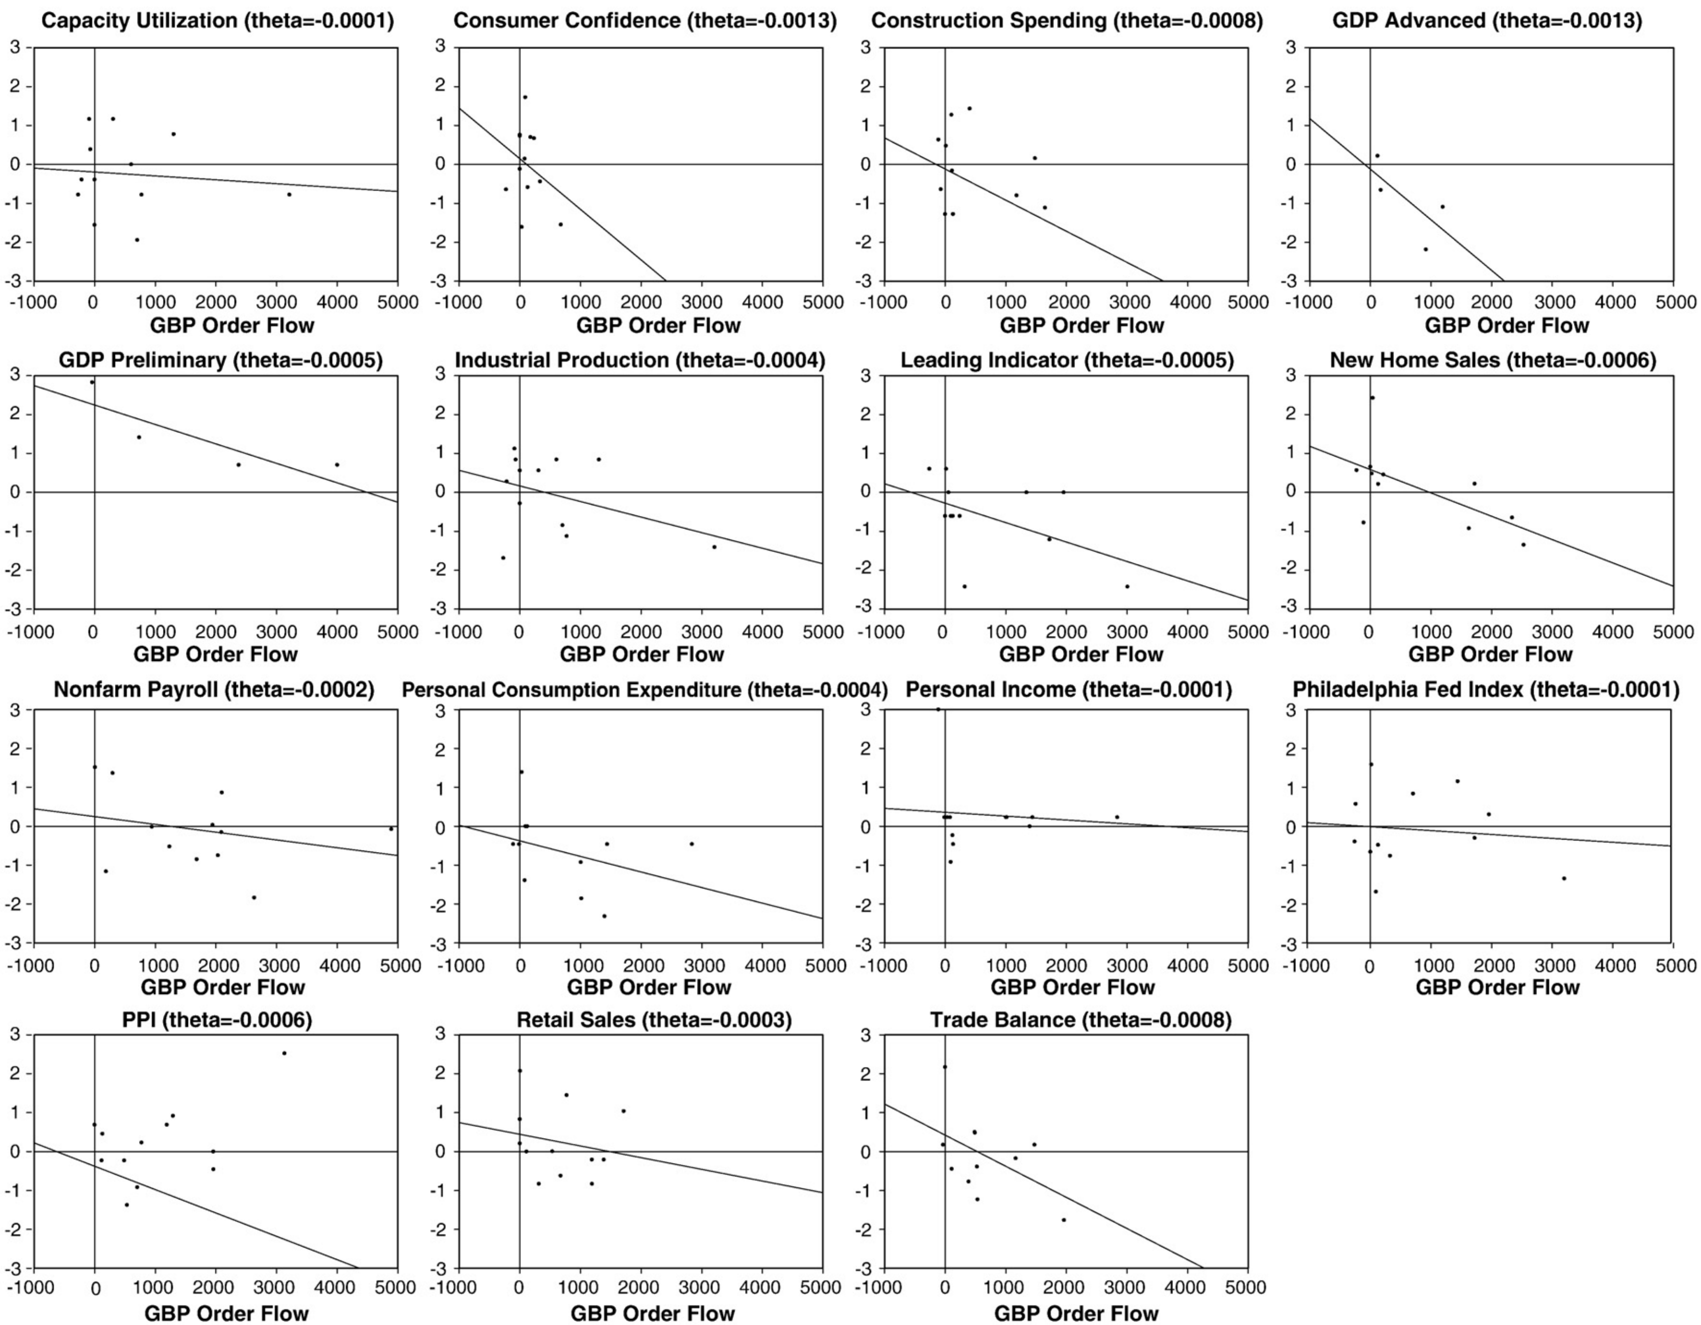

Fig. 4. US news and GBP order flow. The figure presents the scatter plot cumulated GBP order flow (*sumx*, on the horizontal-axis) and standardized expectations gap (on the vertical-axis) for the US news. The line describes the linear relation between order flow and news. Theta is the elasticity of the standardized expectations gap to one unit of order flow.

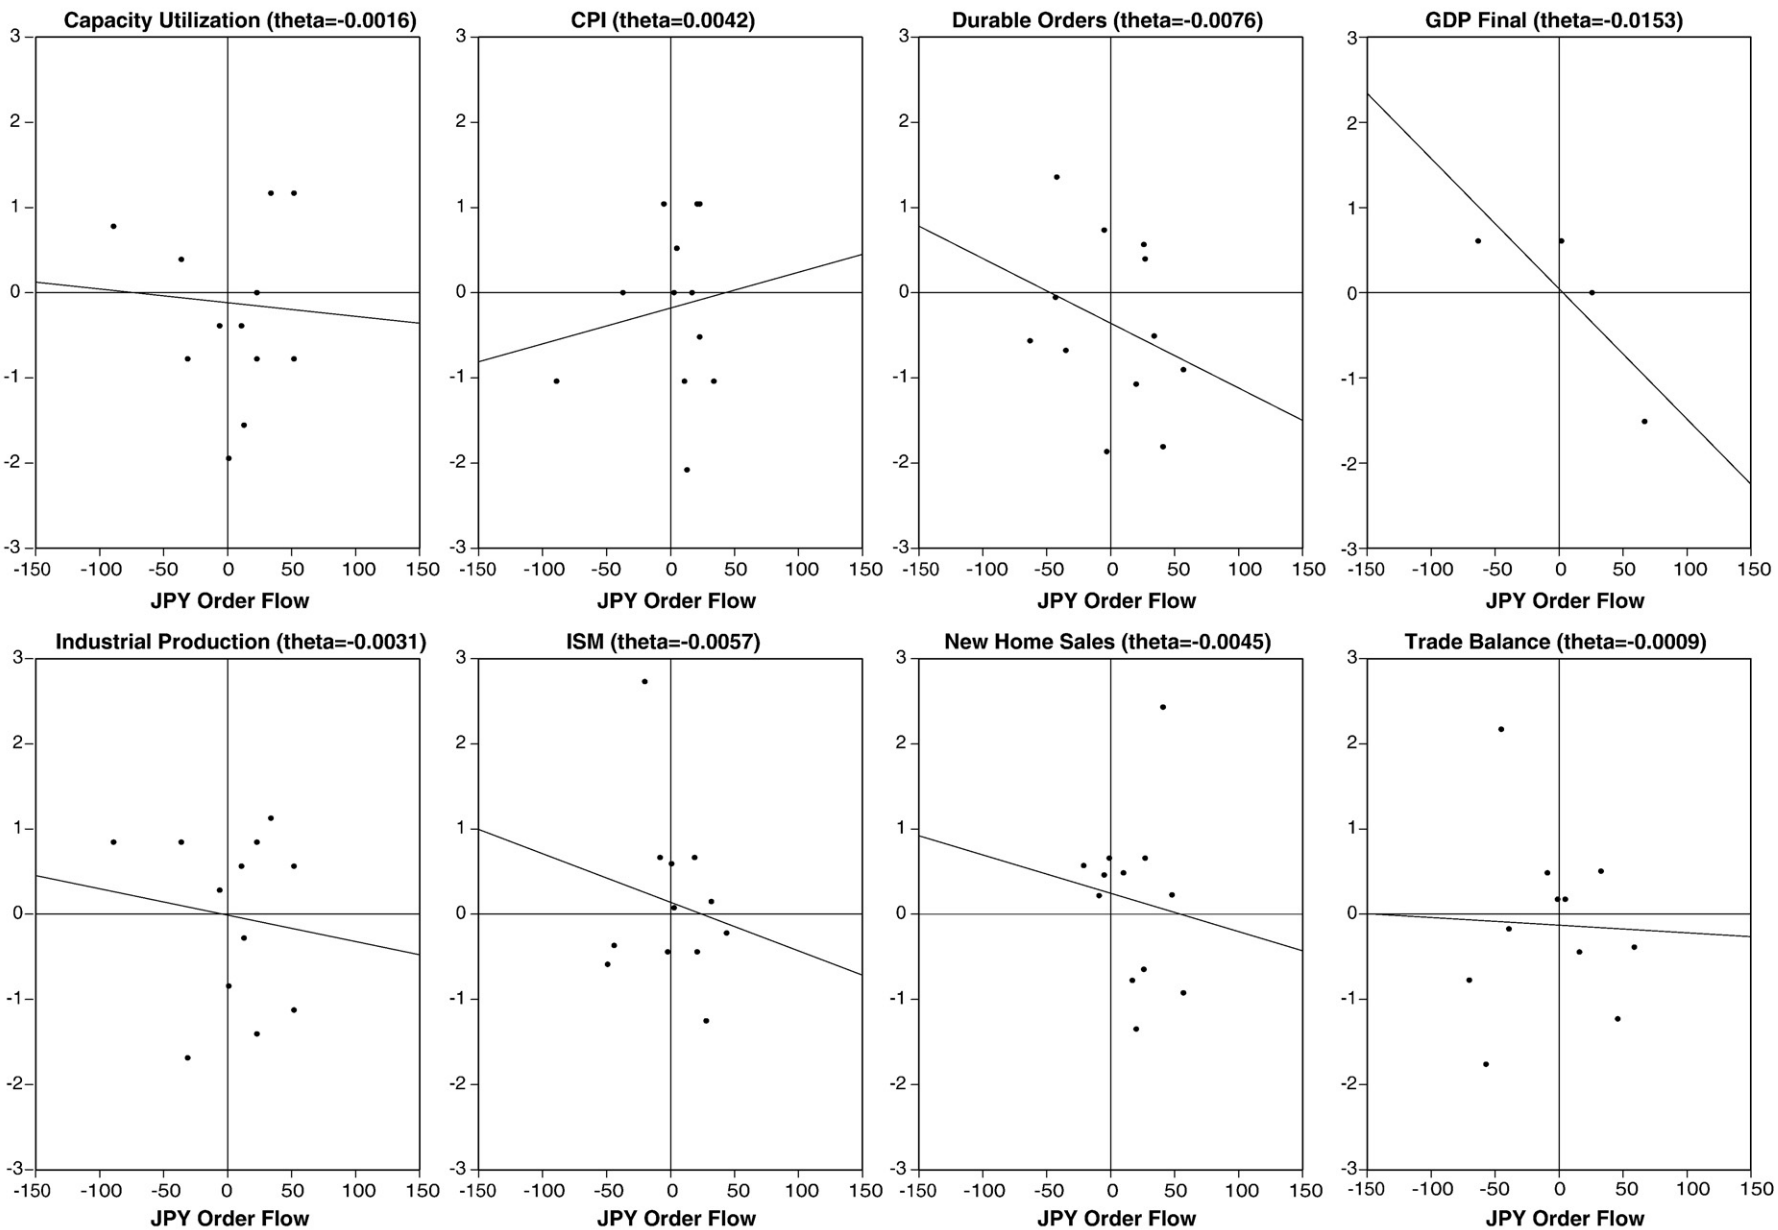

Fig. 5. US news and JPY order flow. The figure presents the scatter plot cumulated JPY order flow (*sumx*, on the horizontal-axis) and standardized expectations gap (on the vertical-axis) for the US news. The line describes the linear relation between order flow and news. Theta is the elasticity of the standardized expectations gap to one unit of order flow.

countries investigated. We take these results as illustrative evidence that supports the conjecture that order flow aggregates changes in market expectations with regard to these fundamentals.[14]

4.3. Summing up

To sum up, the evidence in this section suggests that there is a strong relation between order flow and macroeconomic information. Order flow is intimately linked to both news on fundamentals and to changes in expectations about these fundamentals. Macroeconomic information is identified to be a determinant of order flow, which implies that exchange rate fluctuations may be linked to macroeconomic fundamentals both via a direct link, as in classical exchange rate theory, and via order flow, as in the microstructure approach to FX. These results imply that order flow's explanatory power stems partly from macroeconomic information, lending support to the explanations for the well-documented disappointing results on the direct link between macroeconomic fundamentals and exchange rates in the literature proposed by Evans and Lyons (2002a, 2007, 2008) and Bacchetta and van Wincoop (2006).

5. Empirical models and asset allocation: the framework

Recently, several banks have invested in technology that captures order flow information for forecasting purposes (e.g. the CitiFlow system by Citigroup and similar systems built at UBS, Royal Bank of Scotland and HSBC). The microstructure literature has used some of these data (e.g. Evans and Lyons, 2005a; Marsh and O'Rourke, 2005; Sager and Taylor, 2008) as well as data constructed from electronic platforms, Reuters and EBS (e.g. Evans, 2002; Payne, 2003; Berger et al., 2008). In this section, we examine the forecasting power of order flow in a stylized asset allocation framework, where a mean-variance investor maximizes expected returns subject to a chosen target volatility of portfolio returns.

We rank the performance of the competing models using two main criteria: the Sharpe ratio, arguably the most common measure of performance evaluation among market practitioners; and the performance fee that a risk-averse investor would be willing to pay to switch from a random walk strategy to an active management strategy based on an alternative model. In addition, we calculate the break-even transaction cost, defined as the cost that would remove any economic gain from a dynamic asset allocation strategy relative to a simple random walk strategy. We choose to perform one-day ahead forecasts for the following reasons: one-day ahead forecasts based on order flow are implementable; it is a relevant horizon for practitioners (e.g. most currency funds); unlike intraday forecasts, it involves interest rate considerations; and it is less likely that gradual learning based on this data will allow forecasting at much longer horizons.

5.1. Forecasting models

Consider an investor who forecasts exchange rate returns daily and allocates capital across currencies. We investigate four models that this investor might use: two models based on order flow, a model based on the Fama (1984) regression that exploits the forward bias (carry trade strategy), and a random walk with drift, used as the benchmark against which the other models are evaluated.

A skeptical view of the ability of order flow to explain exchange rates could be that order flow simply captures serial correlations in exchange rate returns rather than genuine fundamentals information. Alternatively, order flow could reflect the impact of exchange rate movements on trading activities via mechanisms of feedback or momentum trading (e.g. Danielsson and Love, 2006). Hence, the first model we consider allows each exchange rate return to depend on lagged order flow of the relevant currency pair, lagged order flow of other currency pairs and also lagged exchange rate changes (capturing momentum) of the currencies examined:

$$\boldsymbol{P}_{t+1} = \boldsymbol{C} + \boldsymbol{\Lambda X}_t + \boldsymbol{\Gamma P}_t + \boldsymbol{U}_{t+1}, \tag{7}$$

where $\boldsymbol{P}_{t+1} = [\Delta s_{t+1}^{\mathrm{EUR}}, \Delta s_{t+1}^{\mathrm{GBP}}, \Delta s_{t+1}^{\mathrm{JPY}}]'$ is the 3×1 vector of exchange rate returns; $\boldsymbol{X}_t$ is the 3×1 vector of order flows; $\boldsymbol{\Lambda}$ and $\boldsymbol{\Gamma}$ are 3×3 matrices of coefficients; $\boldsymbol{C}$ is the vector of constant terms; and $\boldsymbol{U}_{t+1}$ is the vector of error terms. We term the model in Eq. (7) M^{GEN}.

The second model we consider relies only on order flow information, and we term this model 'pure' order-flow model, or M^{POF}. M^{POF} is obtained from imposing $\boldsymbol{\Gamma}=0$ in Eq. (7):

$$\boldsymbol{P}_{t+1} = \boldsymbol{C} + \boldsymbol{\Lambda X}_t + \boldsymbol{U}_{t+1}. \tag{8}$$

The third model examined, M^{FB}, is the well-known 'forward bias' trading strategy based on the Fama (1984) regression:

$$\boldsymbol{P}_{t+1} = \boldsymbol{C} + \boldsymbol{\Pi Z}_t + \boldsymbol{U}_{t+1}, \tag{9}$$

where $\boldsymbol{Z}_t = (1i_t - \boldsymbol{\Upsilon}^*)$ is the 3×1 vector of interest rate differentials (domestic minus foreign); $\boldsymbol{\Upsilon}_t^* = \left[i_t^{\mathrm{EUR}}, i_t^{\mathrm{GBP}}, i_t^{\mathrm{JPY}}\right]'$ denotes the 3×1 vector of foreign interest rates; and $\mathbf{1}$ is a vector of ones.[15]

Finally, the benchmark model is the random walk with drift, M^{RW}:

$$\boldsymbol{P}_{t+1} = \boldsymbol{C} + \boldsymbol{U}_{t+1}. \tag{10}$$

5.2. Asset allocation

This section discusses the framework we use in order to evaluate the impact of predictable changes in exchange rate returns on the performance of dynamic allocation strategies. We employ mean-variance analysis as a standard measure of portfolio performance to calculate Sharpe ratios. Assuming quadratic utility, we also measure how much a risk-averse investor is willing to pay for switching from the naïve random walk strategy that assumes no predictability in exchange rates to a dynamic strategy which conditions on order flow or on the interest rate differential.

5.2.1. Portfolio weights

The investor is assumed to have an initial wealth of $1 million that he invests every day in three risky assets (foreign overnight deposits) and one riskless asset (US overnight deposit). He chooses the weights to invest in each risky asset by constructing a dynamically re-balanced portfolio that maximizes the conditional expected return subject to a target conditional volatility. Let $\mu_{s,t+1|t} = E_t(\boldsymbol{P}_{t+1} + \boldsymbol{\Upsilon}_t^*)$ be the 3×1 vector of conditional expectations for the risky asset returns, then the weights invested in each asset are calculated to solve:

$$\max_{\boldsymbol{w}_t} \left\{ \mu_{p,t+1|t} = \boldsymbol{w}_t' \mu_{s,t+1|t} + \left(1 - \boldsymbol{w}_t' 1\right) i_t \right\} \\ \text{s.t.} \left(\sigma_p^*\right)^2 = \boldsymbol{w}_t' \boldsymbol{\Sigma}_{t+1|t} \boldsymbol{w}_t, \tag{11}$$

[14] In order to circumvent the problems arising from the low number of observations, we also perform the following exercise to examine the relation between cumulative order flow and news. We increase the number of observations by aggregating the news variables in one vector according to different criteria. Then, we estimate a Probit model for the relation between the sign of the cumulative order flow and groups of news: $I_{sumx,t} = \vartheta_0 + \vartheta_1 news_t + \varpi_t$, where $I_{sumx}=1$ if $sumx>0$, and 0 otherwise; *news* is the vector of the grouped news; and ϖ is the error term. The Probit estimation yields a correctly signed and statistically significant coefficient in most cases. For example, in the case of output-related news, better than expected news on output are associated with an increase in demand (hence, the probability of appreciation) for the base currency. These results are not reported to conserve space but are available upon request.

[15] The Fama regression involves the exchange rate return on the left-hand-side and the lagged forward premium on the right-hand-side. Given the use of a daily strategy and the fact that forward contracts for daily maturity do not exist, we replace the forward premium with the overnight interest rate differential, hence relying on the validity of covered interest parity (Akram et al., 2008).

where $\mu_{p,t+1|t}$ is the conditional expected return of the portfolio that combines the risky assets and risk free rate; $\boldsymbol{w}_t$ is the 3×1 vector of portfolio weights on the risky assets; σ_p^* is the target level of risk for the portfolio; $\sum_{t+1|t}$ is the 3×3 variance–covariance matrix of the risky assets and is estimated recursively as the investor updates forecasts and dynamically rebalances his portfolio every day. The solution to this maximization problem yields the risky assets investment weights:

$$\boldsymbol{w}_t = \frac{\sigma_p^*}{\sqrt{Q_t}}\boldsymbol{\Sigma}_{t+1|t}^{-1}\left(\mu_{s,t+1|t} - 1i_t\right), \tag{12}$$

where $\boldsymbol{\mu}_{s,t+1|t} - \mathbf{1}i_t$ is the FX excess return; and $Q_t = (\boldsymbol{\mu}_{s,t+1|t} - \mathbf{1}i_t)'\sum_{t+1|t}^{-1}(\boldsymbol{\mu}_{s,t+1|t} - \mathbf{1}i_t)$. The weight invested in the risk free asset is $1 - \boldsymbol{w}_t'\mathbf{1}$.

The models considered in this paper assume constant volatility (variance–covariance matrix). Hence the only source of time variation in $\sum$ is due to the fact that the models are re-estimated recursively, so that the volatility forecast for time $t+1$ conditioned on information t is equal to the covariance estimated using data up to time t.

5.2.2. Sharpe ratio

The first economic criterion we employ is the Sharpe ratio (SR), or return-to-variability ratio, which measures the risk-adjusted returns from a portfolio or investment strategy and is widely used by investment banks and asset management companies to evaluate investment and trading performance. The ex-post SR is defined as:

$$\mathrm{SR} = \frac{\overline{r_p - r_f}}{\sigma_p}, \tag{13}$$

where $\overline{r_p - r_f}$ is the average (annualized) excess portfolio return over the risk free rate, and σ_p is the (annualized) standard deviation of the investment returns.

This measure is commonly used to evaluate performance in the context of mean-variance analysis. However, Marquering and Verbeek (2004) and Han (2006) show that the SR can underestimate the performance of dynamically managed portfolios. This is because the SR is calculated using the average standard deviation of the realized returns, which overestimates the conditional risk (standard deviation) faced by an investor at each point in time. Thus, we use the performance fee as an additional economic criterion to quantify the economic gains from using the exchange rate models considered.

5.2.3. Performance fees under quadratic utility

We calculate the maximum performance fee a risk-averse investor is willing to pay to switch from the benchmark portfolio (based on the random walk model, M^{RW}) to an alternative portfolio. The specific measure adopted is based on mean-variance analysis with quadratic utility (West et al., 1993; Fleming et al., 2001; Della Corte et al., 2008). Under quadratic utility, at the end of period $t+1$ the investor's utility of wealth can be represented as:

$$U(W_{t+1}) = W_{t+1} - \frac{\varrho}{2}W_{t+1}^2 = W_tR_{p,t+1} - \frac{\varrho}{2}W_t^2R_{p,t+1}^2 \tag{14}$$

where W_{t+1} is the investor's wealth at $t+1$; $R_{p,t+1} = 1 + r_{p,t+1}$ is the gross portfolio return; and ϱ determines his risk preference. To quantify the economic value of each model the degree of relative risk aversion (RRA) of the investor is set to $\delta = \frac{\varrho W_t}{1-\varrho W_t}$, and the same amount of wealth is invested every day. Under these circumstances, West et al. (1993) show that the average realized utility ($\overline{U}$) can be used to consistently estimate the expected utility generated from a given level of initial wealth. The average utility for an investor with initial wealth $W_0 = 1$ is:

$$\overline{U} = \frac{1}{T}\sum_{t=0}^{T-1}\left(R_{p,t+1} - \frac{\delta}{2(1+\delta)}R_{p,t+1}^2\right). \tag{15}$$

Table 4
In-sample performance.

	M^{POF}	M^{GEN}	M^{FB}	M^{RW}
SR	5.79	7.05	2.37	2.23
Φ	29.43	43.76	0.89	–
τ	15.21	14.19	–	–

In-sample performance over the period 2/13/2004–6/14/2004. The dependent variable is Δs_{t+1}, the daily exchange rate return from 17:00 GMT on day t to 17:00 GMT on day $t+1$; daily order flow (positive for net foreign currency purchases) is cumulated between 7:00 and 17:00 GMT, for each exchange rate: EUR, GBP and JPY. Note that only lagged (not contemporaneous) order flow is used in the information set, and the exercise is considered in-sample only because the model parameters are estimated during the period 2/13/2004–6/14/2004 rather than recursively. We present the following in-sample performance criteria: Sharpe ratio (SR), performance fee (Φ, annual percentage points), and break-even transaction cost (τ, daily basis points) for four models: M^{POF}, M^{GEN}, M^{FB} and M^{RW} as described in Section 5.1, for target volatility $\sigma^* = 10\%$. The fees denote the amount an investor with quadratic utility and a degree of relative risk aversion $\delta = 5$ is willing to pay for switching from the random walk benchmark to an alternative model based on other information (order flow and interest rate differential). τ is defined as the minimum daily proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy over the random walk benchmark. "–" indicates that the random walk model is superior to the alternative model. All the evaluation criteria are rounded to the second decimal point.

At any point in time, one set of estimates of the conditional returns is better than a second set if investment decisions based on the first set leads to higher average realized utility, $\overline{U}$. Alternatively, the optimal model requires less wealth to yield a given level of $\overline{U}$ than a suboptimal model. Following Fleming et al. (2001), we measure the economic value of our FX strategies by equating the average utilities for selected pairs of portfolios. Suppose, for example, that holding a portfolio constructed using the optimal weights based on M^{RW} yields the same average utility as holding the optimal portfolio implied by the pure order flow model, M^{POF} that is subject to daily expenses Φ, expressed as a fraction of wealth invested in the portfolio. Since the investor would be indifferent between these two strategies, we interpret Φ as the maximum performance fee he will pay to switch from the M^{RW} to the M^{POF} strategy. In other words, this utility-based criterion measures how much a mean-variance investor is willing to pay for conditioning on order flow as in the M^{POF} strategy for the purpose of forecasting exchange rate returns. The performance fee will depend on the investor's degree of risk aversion. To estimate the fee, we find the value of Φ that satisfies:

$$\sum_{t=0}^{T-1}\left\{\left(R_{p,t+1}^{\mathrm{AM}} - \Phi\right) - \frac{\delta}{2(1+\delta)}\left(R_{p,t+1}^{\mathrm{AM}} - \Phi\right)^2\right\} \tag{16}$$
$$= \sum_{t=0}^{T-1}\left\{R_{p,t+1}^{\mathrm{RW}} - \frac{\delta}{2(1+\delta)}\left(R_{p,t+1}^{\mathrm{RW}}\right)^2\right\},$$

where $R_{p,t+1}^{\mathrm{RW}}$ is the gross portfolio return obtained using forecasts from the benchmark M^{RW} model, and $R_{p,t+1}^{\mathrm{AM}}$ is the gross portfolio return constructed using the forecasts from the alternative model (M^{GEN}, M^{POF} and M^{FB}).

5.2.4. Transaction costs

In dynamic investment strategies, where the individual rebalances the portfolio every day, transaction costs can play a significant role in determining returns and comparative utility gains. However, traders charge transaction costs according to counter-party types and trade size. Thus, instead of assuming a certain cost, we compute the break-even transaction cost τ, which is the minimum daily proportional cost that cancels the utility advantage of a given strategy. We assume that transaction costs at time t equal a fixed proportion τ of the amount traded in currency j:

$$\tau\sum_{j=1}^{3}\left|w_t^j - w_{t-1}^j\left(\frac{1+\Delta s_t^j + i_{t-1}^j}{R_{p,t}}\right)\right|. \tag{17}$$

It is assumed that these costs are the same across currencies, which is consistent with the bid-ask spreads observed in the currency market for EUR, GBP and JPY.

6. The forecasting power of order flow: empirical results

We begin our economic evaluation of one-day-ahead exchange rate predictability by performing in-sample estimations of the four candidate models: M^{GEN}, M^{POF}, M^{FB} and M^{RW}. The estimation is carried out over the period from February 13, 2004 to June 14, 2004, comprising about one third of the available observations. While the number of observations used in the in-sample estimation is relatively small, all models are particularly parsimonious linear models, with a small number of parameters. This allows us to make an assessment of the in-sample performance of the models, to have initial estimates of the parameters over a 'training' period prior to the out-of-sample analysis, and to conduct the latter analysis using two thirds of the observations in the data set.

6.1. Models estimation and in-sample analysis

In our setting, the investor obtains the predicted value of exchange rate returns for 17:00 on day $t+1$, conditioning on order flow information aggregated from 7:00 to 17:00 on day t; he then chooses investment weights and invests in the different currencies using Eqs. (11) and (12). He closes the position at 17:00 on day $t+1$. The in-sample prediction is the fitted value of the exchange rate return for day $t+1$, from 17:00 to 17:00, using the models described in Section 5.1.[16]

The in-sample performance results include the Sharpe ratio, performance fees and break-even transaction costs calculated using Eqs. (13)–(17) for an annual target volatility of $\sigma^*=0.10$ and assuming that the coefficient of relative risk aversion $\delta=5$. The results are presented in Table 4. The Sharpe ratios range between 2.23 and 7.05, pertaining to M^{RW} and M^{GEN} respectively. These Sharpe ratios are very high, but one must keep in mind that these are in-sample calculations over the period from February 13 to June 14, 2004. The best models appear to be M^{POF} and M^{GEN}, which yield Sharpe ratios of 5.79 and 7.05. Also, an investor would be willing to pay large performance fees of 29.43 and 43.76% per annum in order to switch from a random walk strategy to a strategy based on order-flow models M^{POF} and M^{GEN}. The transaction costs that would cancel the above differences in utility between the order flow models and the random walk are above 14 basis points per day.

In short, the two order flow models (M^{POF} and M^{GEN}) deliver fairly similar results, although there is some additional power deriving from lagged exchange rate information (used in M^{GEN} but not in M^{POF}). These results provide *prima facie* evidence of the predictive power of order flow information as compared to two common benchmarks, the forward bias and the random walk models. However, the analysis until this stage is in-sample, while we are ultimately interested in the economic value of order flow as a conditioning variable out-of-sample.

6.2. The out-of-sample economic value of order flow

The in-sample procedure applied so far allowed us to achieve estimation of the relation between spot exchange rates and the conditioning variables used as predictors. In order to assess the usefulness of the exchange rate models, out-of-sample forecasts of spot returns are constructed using all candidate models estimated in the previous sub-section. In particular, we perform one-day-ahead forecasting exercises on the remaining two thirds of the sample, from June 15, 2004 to February 14, 2005. The out-of-sample forecasts are constructed according to a standard recursive procedure, namely conditional only upon information up to the date of the forecast and with successive re-estimation as the date on which forecasts are conditioned moves through the data set, implying that the model parameters are re-estimated daily.

The forecasting results are presented in Panel A of Table 5. The best in-sample model, M^{GEN}, yields a modest Sharpe ratio of 0.44, whereas the best out-of-sample model is M^{POF}. M^{POF} delivers a Sharpe ratio of 1.06, which is very high compared to others found in the literature. For example, Lyons (2001) reports a Sharpe ratio of 0.48 for an equally weighted investment in six currencies.[17] The forward bias and the random walk models exhibit negative Sharpe ratios. This may well be due to the small size of our sample period, since the literature typically records positive risk-adjusted returns from forward bias trading over longer samples (e.g. Della Corte et al., forthcoming). However, the difference in performance between order flow models and the model based on the forward bias has the interesting implication that the forecasting power in order flow stems from fundamentally different information than 'carry trades' of the kind that would be implied by forward bias models. At the same time, the fact that M^{POF} performs better than M^{GEN} also suggests that momentum effects are not particularly important in forecasting exchange rate returns during our sample. This is consistent with recent evidence from Neely et al. (2009) that shows momentum strategies appear to have broken down in the recent years. This finding provides further corroborating evidence that the information in order flow cannot be captured by simple momentum or forward bias strategies, and it is likely to be related to more fundamental information.

Turning to the calculation of performance fees, Panel A of Table 5 presents performance fees for the representative case of an annual target volatility $\sigma^*=0.10$ and relative risk aversion coefficient $\delta=5$. On average an investor is prepared to pay rather high performance fees to switch from the random walk strategy to the order flow based strategies. Specifically, the investor would pay just under 17 annual percentage points of the portfolio value, to switch to the M^{POF} strategy. The performance fee for M^{GEN} is lower (13%) but still sizable. The

Table 5
Out-of-sample performance: one-day ahead.

	M^{POF}	M^{GEN}	M^{FB}	M^{RW}
Panel A. Constant variance				
SR	1.06	0.44	−1.08	−1.27
Φ	16.75	12.54	1.38	–
τ	4.77	4.17	–	–
Panel B. Multivariate GARCH errors				
SR	1.45	0.50	−1.06	−1.35
Φ	24.83	16.28	0.75	–
τ	5.96	4.91	–	–

The table presents out-of-sample performance measures for the four models investigated: M^{POF}, M^{GEN}, M^{FB} and M^{RW} as described in Section 5.1. The out-of-sample period is 6/15/2004–2/14/2005. The investment is based on the one-period-ahead forecasts generated by the models and calculated from the investment strategy detailed in Section 5.2, for an annual target volatility $\sigma^*=10\%$. We present: the Sharpe ratio (SR), performance fee (Φ, annual percentage points), and the break-even transaction cost (τ, daily basis points). The fees denote the amount an investor with quadratic utility and a degree of relative risk aversion $\delta=5$ is willing to pay for switching from the random walk benchmark to an alternative model based on other information (order flow and interest rate differential). τ is defined as the minimum daily proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy over the random walk benchmark. "–" indicates that the random walk model is superior to the alternative model. All the evaluation criteria are rounded to the second decimal point. In Panel A we report results from models where the error term is assumed to have a constant variance. In Panel B we report the results for the case where the error terms are modeled as a multivariate GARCH process.

[16] Put another way, the investor forecasts the change in the exchange rate between 17:00 of day t and 17:00 of day $t+1$ and closes the position at 17:00 on day $t+1$, realizing a log return $\Delta S_{t+1}=S_{t+1}^{17:00}-S_t^{17:00}$.

[17] For equities, the typical Sharpe ratio from a buy-and-hold strategy in the S&P500 is around 0.4 (Sharpe, 1994; Lyons, 2001; Sarno, 2005). Research on fund performance shows that hedge funds achieve average SRs of 0.36 for the period 1988–1995 (Ackermann et al., 1999), while the average SRs for off-shore hedge funds range from 0.94 to 1.19 (Brown et al., 1999).

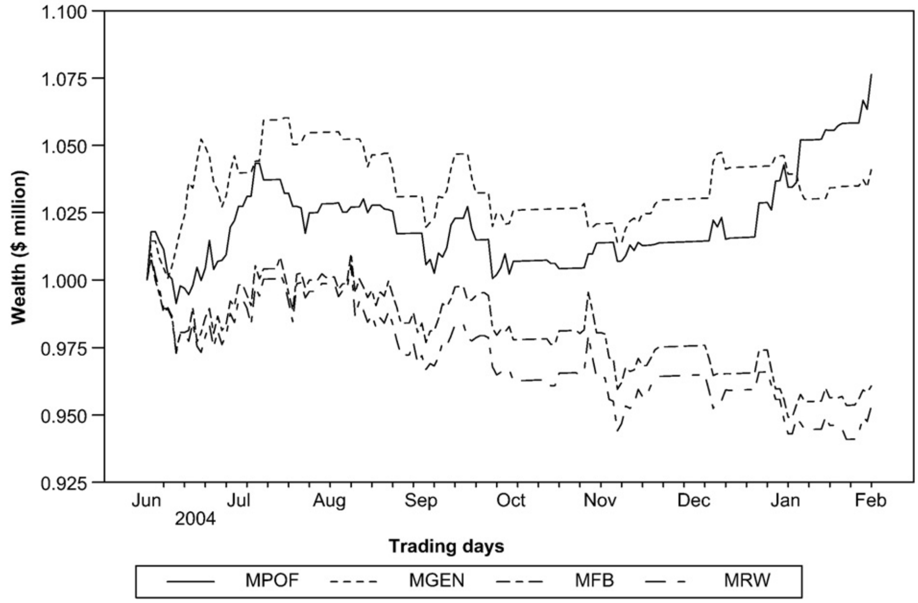

Fig. 6. Wealth evolution out-of-sample. The out-of-sample wealth evolution for each trading hour, for the period 6/15/2004–2/14/2005. The solid line presents the pure order flow M^{POF} wealth evolution, while the other dashed lines present the other models, as indicated in the legend. All investments start from an equal initial wealth of $1 million.

investor would also be willing to shift from the random walk model to the forward bias model, but the maximum performance fee he would be willing to pay is very small. The results are consistent with what was previously found using the Sharpe ratio criterion and are quite high compared to those found in the literature.[18]

Panel A of Table 5 also presents the break-even transaction costs (τ). The results show that conditioning on order flow information leads to high τ values. Specifically, the break-even transaction cost for model M^{POF} is 4.77 basis points per day. This break-even transaction cost is much higher than the transaction costs implied by the bid-ask spread observed in the market, making it highly unlikely that the transaction costs would cancel the high returns generated from M^{POF}. The corresponding τ value for M^{GEN} is 4.17, also comfortably high to conclude that the economic value of order flow allows to build strategies that dominate the random walk and forward bias benchmarks even after allowing for reasonable transactions costs—the bid-ask spread for the exchange rates examined here is in the range between 1 and 2 basis points.

To provide a visual illustration of the portfolio results obtained using the four candidate models, we present in Fig. 6 the evolution of wealth for M^{POF}, M^{GEN}, M^{FB} and M^{RW}. From this graph, we notice that the high Sharpe ratios under M^{POF} and, to a lesser extent, M^{GEN} are due to relatively high returns for the desired target investment volatility. In fact, the wealth evolution under M^{POF} and M^{GEN} appears very close for a long period of time where M^{GEN} performs slightly better, and the outperformance of M^{POF} over M^{GEN} is due to losses incurred under M^{GEN} towards the end of the sample.[19]

Panel B of Table 5 reports out-of-sample results for the case where we depart from the assumption of a constant variance in the error term. As documented in West et al. (1993), Fleming et al. (2001) and Della Corte et al. (forthcoming), adequate modeling of conditional volatility provides economic value in an asset allocation context. For illustrative and robustness purposes, we carry out the out-of-sample asset allocation problem with the errors assumed to follow a multivariate GARCH (1,1) process—the estimation results from the multivariate GARCH(1,1) are not reported to conserve space. The outcome from this exercise confirms that the performance fees, Sharpe ratios and the break-even transaction costs increase when modeling conditional volatility.

Finally, we investigate the forecasting power of order flow in the context of an asset allocation strategy for horizons longer than one day ahead. This exercise provides evidence on the speed at which the forecasting power of order flow information decays over time. If collecting real-time order flow information were too cumbersome, for example, and the investor had to act on day t on order flow information up to day $t-1$ for forecasting exchange rate returns on day $t+1$, effectively the model would be successful only if the information

Table 6
Out-of-sample performance: two- and five-day ahead.

	M^{POF}	M^{GEN}	M^{FB}	M^{RW}
Panel A. Two-day ahead				
SR	0.94	0.39	−1.08	−1.27
Φ	16.00	12.21	1.38	–
τ	4.59	4.10	–	–
Panel B. Five-day ahead				
SR	1.14	0.61	−1.08	−1.27
Φ	17.23	13.50	1.38	–
τ	5.40	4.98	–	–

The table presents out-of-sample performance measures at the two- and five-day horizon for the four models investigated: M^{POF}, M^{GEN}, M^{FB} and M^{RW} as described in Section 5.1. The out-of-sample period is 6/15/2004–2/14/2005. The investment is based on the forecasts generated by the models and calculated from the investment strategy detailed in Section 5.2, for an annual target volatility $\sigma^* = 10\%$. We present: the Sharpe ratio (SR), performance fee (Φ, annual percentage points), and the break-even transaction cost (τ, daily basis points). The fees denote the amount an investor with quadratic utility and a degree of relative risk aversion $\delta = 5$ is willing to pay for switching from the random walk benchmark to an alternative model based on other information (order flow and interest rate differential). τ is defined as the minimum daily proportional cost which cancels out the utility advantage (and hence positive performance fee) of a given strategy over the random walk benchmark. "–" indicates that the random walk model is superior to the alternative model. We report results from all models assuming that the error term has a constant variance. All the evaluation criteria are rounded to the second decimal point.

[18] The difference between the realized portfolio returns of M^{POF} and the corresponding realized returns of the random walk benchmark is also statistically significant at the 5% significance level, with a t-statistic of 1.86. This is not the case for M^{GEN} and M^{FB}, for which we find t-statistics of 1.19 and 0.48, respectively.

[19] These results may be considered conservative because in our setting the investor is forced to close his position at a specific time of the day (17:00); realistically, he could place limit orders that allow him to make higher profits in the day. The investor can place a limit order to sell at an exchange rate level higher than the forecast exchange rate, and if the order is filled he makes even higher profits. Furthermore, a trader can place limit orders that expire at every hour of the day, or a combination of stop-loss and take-profit orders, thus accumulating more profits during the day.

in order flow can forecast exchange rate returns up to two days ahead. In essence, this exercise sheds light on the robustness of our previous results with respect to lags in the available information set for order flow. We consider both two- and five-day (one week) ahead forecasts. Note that we do not lag the information set in the case of M^{FB} and M^{RW} because they are based on readily available information. The results, reported in Table 6, indicate that the ranking of models in terms of performance measures remains the same as in Table 5, with M^{POF} and, to a lesser extent, M^{GEN} outperforming M^{FB} and M^{RW} both in two- and five-day ahead forecasting.

7. Conclusions

This paper makes two related contributions to empirical exchange rate economics. We show that order flow is related to current and expected future macroeconomic fundamentals, and can profitably forecast risk-adjusted currency returns.

Previous research has found that order flow has strong explanatory power for exchange rate movements, whereas macroeconomic fundamentals have weak explanatory power. We provide evidence that a significant amount of order flow variation can be explained using macroeconomic news, suitably constructed from survey data. In addition, order flow appears to aggregate changes in expectations about fundamentals. This finding may provide a rationale for the high explanatory power of order flow found in the literature and complements the evidence that macro information affects exchange rates at high frequency (Andersen et al., 2003) and forecasts exchange rate returns at long horizons (e.g. Mark, 1995). Furthermore, this result suggests that the order flow channel is key to link exchange rates to fundamentals, as argued by Evans and Lyons (2002a, 2007) and Bacchetta and van Wincoop (2006).

The well-documented inability of standard exchange rate models to forecast out-of-sample better than a naïve random walk has been and perhaps remains the conventional wisdom in the international finance profession. However, if exchange rates are determined by macroeconomic fundamentals, but order flow gradually conveys information on heterogeneous beliefs about these fundamentals, then order flow should provide forecasting power for exchange rates. The key finding of this paper is that order flow provides powerful information that allows us to forecast the daily exchange rate movements of three major exchange rates. This result is obtained by measuring forecasting power in the context of simple, intuitive metrics of economic gains. We show the Sharpe ratios and the performance fees against the random walk strategy for a mean-variance investor that uses out-of-sample exchange rate forecasts obtained from different models that condition on order flow or forward bias. The Sharpe ratio from using the order flow model is well above unity and substantially higher than any alternative model considered, while the performance fees are just under 17% per annum.

In summary, taking together the results provided in this paper, we add further evidence that order flow is crucial to understanding exchange rate fluctuations. Order flow is strongly related to fundamentals and, in turn, can provide useful guidance to forecast exchange rate movements. We also take this as evidence that can bridge the micro–macro divide, in the sense that current and future exchange rates are not random walks but are, at least indirectly, determined by economic fundamentals.

While these results aid the profession's understanding on the behavior of exchange rates and the connection between the state of the economy and currency trading activities, we view our results only as one step ahead in the relevant debate. Macroeconomic information, order flow and exchange rates are linked by complex dynamic interactions, and much more work needs to be carried out to shed light on this relation. It would seem fruitful to investigate currencies other than the major and more liquid dollar exchange rates analyzed here. It would also seem logical to extend this line of research to the analysis of the link between order flow and exchange rate volatility, which is notoriously difficult to explain in terms of plain macro variables in a similar vain as exchange rate returns (Berger et al., 2009). These issues remain on the agenda for future research.

Acknowledgments

This paper was partly written while Lucio Sarno and Elvira Sojli were visiting Norges Bank. We are grateful for constructive comments to Charles Engel (editor), an anonymous referee, Tam Bayoumi, Geir Bjønnes, Hilde Bjørnland, Alain Chaboud, Pasquale Della Corte, Frank Diebold, Martin Evans, Bilal Hafeez, Robert Kosowski, Richard Lyons, Ian Marsh, Lukas Menkhoff, Michael Moore, Carol Osler, Anna Pavlova, Ilias Tsiakas, Bent Vale, Giorgio Valente, Eric van Wincoop, Paolo Vitale, to other participants to the ESRC-WEFRP/IMF International Macro-Finance Conference, 2008; Western Finance Association, 2008; Royal Economic Society Conference, 2008; European Finance Association, 2008; Far Eastern Meeting of the Econometric Society, 2008; INFINITI Conference, 2008; European Economic Association, 2007; Annual Conference on "Advances in the Analysis of Hedge Fund Strategies," Imperial College Business School, 2007; Microstructure of Foreign Exchange and Equity Markets Workshop, Bank of Canada, 2006; and to participants at seminars held at the Rotterdam School of Management, University of Copenhagen, Manchester Business School, Universidad Carlos III, Federal Reserve Bank of St. Louis, Bank of Canada, Norges Bank, and Warwick Business School. The authors alone are responsible for the views expressed in the paper and for any errors that may remain.

References

Abhyankar, A., Sarno, L., Valente, G., 2005. Exchange rates and fundamentals: evidence on the economic value of predictability. J. Int. Econ. 66, 325–348.
Ackermann, C., McEnally, R., Ravenscraft, D., 1999. The performance of hedge funds: risk, return, and incentives. J. Finance 54, 833–874.
Admati, A.R., Pfleiderer, P., 1988. A theory of intraday patterns: volume and price variability. Rev. Financ. Stud. 1, 3–40.
Akram, Q.F., Rime, D., Sarno, L., 2008. Arbitrage in the foreign exchange market: turning on the microscope. J. Int. Econ. 76, 237–253.
Andersen, T.G., Bollerslev, T., Diebold, F.X., Vega, C., 2003. Micro effects of macro announcements: real-time price discovery in foreign exchange. Am. Econ. Rev. 93, 38–62.
Bacchetta, P., van Wincoop, E., 2004. A scapegoat model of exchange-rate fluctuations. Am. Econ. Rev. Pap. Proc. 94, 114–118.
Bacchetta, P., van Wincoop, E., 2006. Can information heterogeneity explain the exchange rate determination puzzle? Am. Econ. Rev. 96, 552–576.
Bank for International Settlements, 2005. Triennial Central Bank Survey of Foreign Exchange and Derivative Market Activity in 2004. Bank for International Settlements, Basel.
Berger, D.W., Chaboud, A.P., Chernenko, S.V., Howorka, E., Wright, J.H., 2008. Order flow and exchange rate dynamics in electronic brokerage system data. J. Int. Econ. 75, 93–109.
Berger, D.W., Chaboud, A.P., Hjalmarsson, E., 2009. What drives volatility persistence in the foreign exchange markets? J. Financ. Econ., forthcoming.
Bjønnes, G.H., Rime, D., 2005. Dealer behavior and trading systems in foreign exchange markets. J. Financ. Econ. 75, 571–605.
Bjønnes, G.H., Rime, D., Solheim, H.O.A., 2005. Liquidity provision in the overnight foreign exchange market. J. Int. Money Financ. 24, 177–198.
Breedon, F., Vitale, P., 2005. An empirical study of liquidity and information effects of order flow on exchange rates. mimeo, Imperial College Business School.
Brown, S.J., Goetzmann, W.N., Ibbotson, R.G., 1999. Offshore hedge funds: Survival and performance, 1989–95. J. Bus. 72, 91–117.
Chakravarty, S., 2001. Stealth-trading: which traders' trades move stock prices? J. Financ. Econ. 61, 289–307.
Cheung, Y.W., Chinn, M.D., Garcia Pascual, A., 2005. Empirical exchange rate models of the nineties: are any fit to survive? J. Int. Money Financ. 24, 1150–1175.
Chinn, M.D., Moore, M.J., 2008. Private information and a macro model of exchange rates: Evidence from a novel data set. NBER Working Paper 14175.
Daníelsson, J., Love, R., 2006. Feedback trading. Int. J. Financ. Econ. 11, 35–53.
Daníelsson, J., Luo, J., Payne, R., 2002. Exchange rate determination and inter-market order flow effects. typescript, London School of Economics.
Della Corte, P., Sarno, L., Thornton, D.L., 2008. The expectation hypothesis of the term structure of very short-term rates: statistical tests and economic value. J. Financ. Econ. 89, 158–174.
Della Corte, P., Sarno, L., Tsiakas, I., forthcoming. An economic evaluation of empirical exchange rate models. Review of Financial Studies. doi:10.1093/rfs/hhn058.
Dominguez, K.M., Panthaki, F., 2006. What defines 'News' in foreign exchange markets? J. Int. Money Financ. 25, 168–198.

Easley, D., O'Hara, M., 1992. Adverse selection and large trade volume: the implications for market efficiency. J. Financ. Quant. Anal. 27, 185–208.
Elliott, G., Ito, T., 1999. Heterogeneous expectations and tests of efficiency in the yen/dollar forward exchange rate market. J. Monet. Econ. 43, 435–456.
Engel, C., West, K.D., 2005. Exchange rates and fundamentals. J. Polit. Econ. 113, 485–517.
Engel, C., Mark, N.C., West, K.D., 2008. Exchange rate models are not as bad as you think. In: Acemoglu, D., Rogoff, K., Woodford, M. (Eds.), NBER Macroeconomics Annual 2007. MIT Press, Cambridge, MA, pp. 381–441.
Evans, M.D.D., 2002. FX trading and exchange rate dynamics. J. Finance 57, 2405–2447.
Evans, M.D.D., Lyons, R.K., 2002a. Order flow and exchange rate dynamics. J. Polit. Econ. 110, 170–180.
Evans, M.D.D., Lyons, R.K., 2002b. Informational integration and FX trading. J. Int. Money Financ. 21, 807–831.
Evans, M.D.D., Lyons, R.K., 2005a. Meese–Rogoff redux: micro-based exchange-rate forecasting. Am. Econ. Rev. Pap. Proc. 95, 405–414.
Evans, M.D.D., Lyons, R.K., 2005b. Do currency markets absorb news quickly? J. Int. Money Financ. 24, 197–217.
Evans, M.D.D., Lyons, R.K., 2006. Understanding order flow. Int. J. Financ. Econ. 11, 3–23.
Evans, M.D.D., Lyons, R.K., 2007. Exchange rate fundamentals and order flow. NBER Working Paper 13151.
Evans, M.D.D., Lyons, R.K., 2008. How is macro news transmitted to exchange rates? J. Financ. Econ. 88, 26–50.
Fama, E.F., 1984. Forward and spot exchange rates. J. Monet. Econ. 14, 319–338.
Faust, J., Rogers, J.H., Wright, J.H., 2003. Exchange rate forecasting: the errors we've really made. J. Int. Econ. 60, 35–59.
Fleming, J., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing. J. Finance 56, 329–352.
Galati, G., Melvin, M., 2004. Why has FX trading surged? Explaining the 2004 triennial survey. BIS Q. Rev. 67–74.
Glosten, L.R., Milgrom, P.R., 1985. Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. J. Financ. Econ. 14, 71–100.
Han, Y., 2006. Asset allocation with a high dimensional latent factor stochastic volatility model. Rev. Financ. Stud. 19, 237–271.
Killeen, W.P., Lyons, R.K., Moore, M.J., 2006. Fixed versus flexible: lessons from EMS order flow. J. Int. Money Financ. 25, 551–579.
Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335.
Leitch, G., Tanner, J.E., 1991. Economic forecast evaluation: Profits versus the conventional measures. Am. Econ. Rev. 81, 580–590.
Love, R., Payne, R., 2008. Macroeconomic news, order flows and exchange rates. J. Financ. Quant. Anal. 43, 467–488.
Lyons, R.K., 2001. The Microstructure Approach to Exchange Rates. MIT Press, Cambridge, MA.
Mark, N.C., 1995. Exchange rates and fundamentals: evidence on long-horizon predictability. Am. Econ. Rev. 85, 201–218.
Marquering, W., Verbeek, M., 2004. The economic value of predicting stock index returns and volatility. J. Financ. Quant. Anal. 39, 407–429.
Marsh, I.W., O'Rourke, C., 2005. Customer order flow and exchange rate movements: is there really information content? Working paper, Cass Business School.
Meese, R.A., Rogoff, K., 1983. Empirical exchange rate models of the seventies: do they fit out of sample? J. Int. Econ. 14, 3–24.
Molodtsova, T., Papell, D.H., 2009. Out-of-sample exchange rate predictability with Taylor rule fundamentals. J. Int. Econ. 77, 167–180.
Neely, C.J., Weller, P.A., Ulrich, J.M., 2009. The adaptive market hypothesis: evidence from the foreign exchange market. J. Financ. Quant. Anal., forthcoming.
O'Hara, M., 1995. Market Microstructure Theory. Blackwell, Cambridge, MA.
Payne, R., 2003. Informed trade in spot foreign exchange markets: an empirical investigation. J. Int. Econ. 61, 307–329.
Rime, D., 2003. New electronic trading systems in the foreign exchange markets. In: Jones, D.C. (Ed.), New Economy Handbook. Academic Press, San Diego.
Rossi, B., 2006. Are exchange rates really random walks? Some evidence robust to parameter instability. Macroecon. Dyn. 10, 20–38.
Sager, M.J., Taylor, M.P., 2008. Commercially available order flow data and exchange rate movements: caveat emptor. J. Money, Credit Bank. 40, 583–625.
Sarno, L., 2005. Towards a solution to the puzzles in exchange rate economics: where do we stand? Can. J. Econ. 38, 673–708.
Sarno, L., Valente, G., 2009. Exchange rates and fundamentals: footloose or evolving relationship? Journal of the European Economic Association, forthcoming.
Sharpe, W.F., 1994. The Sharpe ratio. J. Portf. Manage. 21, 49–58.
West, K.D., Edison, H.J., Cho, D., 1993. A utility based comparison of some models for exchange rate volatility. J. Int. Econ. 35, 23–45.