CUSTOMER ORDER FLOW AND EXCHANGE RATE MOVEMENTS: IS THERE REALLY INFORMATION CONTENT?

Ian W. Marsh and Ceire O'Rourke*

First Draft: April 2004

This Draft: April 2005

Abstract:

In this paper we analyse the information content of the customer order flow seen by a leading European commercial bank's foreign exchange desk. We attempt to distinguish between three different explanations given in the literature for the positive contemporaneous correlation between exchange rate changes and net order flows. We discount the liquidity effect since otherwise equivalent order flows from different counterparties have different correlations with exchange rate changes. While it is harder to discount the feedback trading explanation we find evidence that a measure of the degree of informedness of customers widely used in the equity microstructure literature closely corresponds to the size of the correlation between order flow and exchange rate changes. We argue that customer order flows do contain information.

Keywords: Foreign exchange, customer order flow, PIN

* Cass Business School, London. The authors would like to than Jakob Lage Hansen, Soeren Hvidkjaer, Roberto Rigobon, Paulo Vitale and seminar participants at the Bank of England, Danish National Bank and Cass Business School for comments. We are particularly grateful to the Royal Bank of Scotland for providing the order flow data and for discussing the realities of the foreign exchange market place with us.

A large body of literature beginning with Mark (1985) suggests that macroeconomic fundamentals (money supplies, prices and income levels) can explain exchange rate movements over horizons in excess of two years. Isolated papers claim that macroeconomic models can provide acceptable forecasting power over horizons as short as three-months (e.g. MacDonald and Marsh, 2004). However, decades of academic research in the field of macroeconomic exchange rate behaviour has failed to provide a convincing explanation of short-term currency movements. The foreign exchange microstructure literature, inspired by Lyons (1995), is still small but has something more positive to say about short-horizon exchange rate movements. The most promising results so far are that there is a positive correlation between spot exchange rate movements and order flows in the inter-dealer market (Evans and Lyons, 2002a) and between spot exchange rate movements and customer order flows (Fan and Lyons, 2003).

The cause of these correlations is not clear. Three are often suggested, two of which are based on causation running from flows to exchange rates. First, there may be private information contained in customer order flow (and reflected in inter-dealer order flows) that is relevant for the valuation of a currency in a non-transitory way. This information may be related to the payoffs from holding the currency (e.g. future interest rates) or to the discount rates that should be applied to future payoffs. Portfolio balance effects are one explanation for persistent time-variation in discount rates that are closely related to order flows. A second explanation for the correlation is that there are transitory liquidity effects on exchange rates caused, for example, by inventory considerations in the pricing behaviour of foreign exchange dealers. Dealers charge a temporary risk premium to absorb unwanted inventory that affects the exchange rate only for as long as the dealer community has to hold the unwanted inventory. Since risk sharing is rapid in the foreign exchange market, these liquidity effects on the exchange rate are likely to be transitory. The final explanation for the correlation reverses the causality and argues that changes in the exchange rate induce flows – so-called feedback trading. The positive correlation could be due to customers buying (selling) a currency that has just appreciated (depreciated).

This paper attempts to differentiate between explanations by considering evidence from a new daily data set of customer order flows covering almost two years, provided by a leading European commercial bank. Compared with data sets used so far it has several advantages. First, the data used in this paper are from six bilateral exchange rates between four currencies (euro, dollar, yen and pound). This allows for a more comprehensive analysis of inter-currency information content than is usually performed since other customer order data sets are limited to one or two exchange rates. Second, the order flow data are broken down according to the nature of the customer, allowing us to test for different information content according to customer type. Third, the data give the value of the order flows rather than simply the number of buys and sells. In analysing customer orders, which are highly non-standard in value, unlike the inter-dealer market, this is an essential characteristic of the data. Finally, the data give gross order flows (buy and sell) rather than simply net order flows (buys minus sells). This allows us to apply, for the first time in the foreign exchange market literature, an estimation procedure based on net and gross order flows that has proved useful in distinguishing between information content of trades in equity markets.

Our main results include the following. First, we confirm that order flows – in our case *customer* order flows – are associated with contemporaneous exchange rate movements at both the daily and weekly frequency. Second, and again consistent with previous research, we find that different components of the order flow have different correlations with exchange rate movements. In particular, order flows from non-financial corporate customers are negatively correlated with exchange rate changes, while flows from financial companies are positively correlated with exchange rate movements. This suggests that liquidity effects are not behind the correlation since if they were, otherwise equivalent order flow from different customer classes should impact the exchange rate equally. However, the negative correlation between exchange rates and non-financial corporate order flows is hard to justify in an information-related framework. Third, and as far as we know new to the literature, we show that information relevant to one exchange rate is contained in customer order flows observed for other exchange rates.[1] Finally we show that the correlation between exchange rate changes and customer order flow is itself highly

[1] Evans and Lyons (2002c) report similar findings from a system of nine bilateral rates against the US dollar using direct inter-dealer order flows. Danielsson, Payne and Luo (2002) find cross-market effects using brokered inter-dealer flows.

positively correlated with the probability of information-based trading measure developed by Easley, Keifer and O'Hara (1996, 1997a, 1997b). This, we argue, is further evidence that (financial) customer order flows contain information relevant for exchange rate determination.

The rest of the paper is arranged as follows. Section 1 surveys the existing research in foreign exchange microstructure with an emphasis on the role of customer orders flows. Section 2 contains a discussion of the data set. Since the data is highly confidential this is rather short. Section 3 presents the results of a series of regressions of changes in exchange rates on customer order flow data and section 4 estimates the probability of information-based trading model. Section 5 shows the degree of consistency between the key results of the previous two sections, and section 6 concludes.

1. Customer order flows and exchange rate movements

Inter-dealer order flow has been the empirical focus of the foreign exchange microstructure literature, primarily because of a lack of data on customer order flows. Inter-dealer order flow data, either from direct inter-dealer trading platforms (Evans and Lyons, 2002a) or the broker platforms (Payne, 2003), have been shown to be highly correlated with changes in spot exchange rates. Coefficients of determination in excess of 0.6 from a regression of spot rate changes on daily signed order flow have spurred interest in microstructure, especially given the awful performance of macro approaches to exchange rates.

However, it is a stylised fact that foreign exchange dealers open and close their trading day with zero inventory positions (Lyons, 1998; Bjønnes and Rime, 2003). The impetus for dealers to trade often comes from orders initiated by their (non-dealer) customers. For example, suppose customer 1 sells €5m to dealer A in exchange for US dollars. Dealer A now has a positive euro (negative dollar) inventory that needs to be managed. That inventory can be reduced in two ways.[2]

[2] We do not discuss a third alternative, that the dealer could hedge his exposure using another instrument (e.g. options) since Fan and Lyons (2003) argue that this is very rare in foreign exchange markets.

First, Dealer A can sell euros to another customer in exchange for dollars. Customer trades are always instigated by the customer, but dealer A can attract customers by offering advantageous rates – so-called price shading. Second, the dealer can pass inventory on to other dealers, either directly or via brokers. The inventory that enters the inter-dealer network then becomes a "hot potato" (Lyons, 1997) and is passed from dealer to dealer. It exits the inter-dealer network when offset by one or more customer orders to buy euros. Figure 1 illustrates this process. If all dealers successfully target a zero inventory position, then, by definition, the sum of all signed customer order flow must also be zero. However, the sum of all signed inter-dealer order flow will usually be non-zero. Inter-dealer order flows magnify the initial customer order depending on how many times the order is passed on and how much leakage to customers occurs in the process.[3] But the whole process is initiated by customer orders, and ended by customer orders (taking, in aggregate, the opposite side of the initial customer order).

Aggregated over a trading day, total signed customer orders should be zero, at least as an approximation.[4] Therefore, daily customer order flow from a representative bank should be only randomly different from zero and uncorrelated with exchange rate changes. However, individual banks may not be representative of the market as a whole. Fan and Lyons (2003) argue that some banks may have disproportionately high shares of what they call "high-impact" customers. They find support for this alternative since cumulative customer order flow from Citibank is highly correlated with exchange rate movements. One explanation of the higher than average impact of Citibank customers could be that they are, on average, better informed. The transactions of Citibank's customers partially reveal this information to Citibank's

---

[3] This assumes inventory is passed on through aggressive strategies such as instigating a direct inter-dealer transaction or by placing market orders in broker systems. Passive strategies such as posting a limit-order would not magnify the initial order. Trading on the basis of a customer's order would, of course, increase the magnification.

[4] This is only an approximation since the foreign exchange market never closes. Throughout any 24-hour weekday period there is an open dealer community capable of holding inventory from customers. However, some periods of the 24-hour window are relatively thin (specifically after the US closes and before London opens) and dealers active then are not capable of carrying a large inventory. Further, all significant markets close over weekends.

dealers, and the subsequent actions of Citibank's dealers (be that price shading to attract other customer orders or transactions in the inter-dealer market) partly reveal Citibank's information advantage. The market slowly learns from customer and inter-dealer transactions and the information is impounded in the spot price. This informational interpretation lies at the heart of much work on order flow.[5]

However, alternative interpretations exist. First, customers that quickly buy a currency they have just observed appreciate would lead to positive correlation between the exchange rate and net order flows at the daily frequency. The information approach assumes the observed correlation is due to causality from trades to exchange rates, but positive feedback trading reverses the direction of causality (Danielsson and Love, 2004).[6] Intraday data on foreign exchange order flow suggests that, if anything, there is negative feedback in the inter-dealer market (Evans and Lyons, 2002b) but to our knowledge there has been no work done on feedback trading in customer flows.

Second, Evans and Lyons (2002a) suggest that risk-averse dealers need to be compensated for absorbing customer order flows by a shift in the exchange rate. In this case, causation running from order flow to the exchange rate leads to the positive correlation but it has nothing to do with information content and is instead due to illiquidity in the market. Breedon and Vitale (2004) model this formally and present evidence, based unfortunately on brokered inter-dealer flows, suggesting that inventory effects account for almost all of the effect of order flow.[7]

Analyses of customer order flows are rare, primarily because banks are understandably reluctant to divulge such sensitive information. Lyon's work noted

[5] Foreign exchange dealers themselves also believe that access to a large customer base conveys a competitive advantage (Cheung, Chinn and Marsh, 2000).

[6] Using tick-by-tick data, Cohen and Shin (2003) provide evidence that price declines (increases) elicit sales (purchases) in the US Treasury note market, particularly during periods of market stress. The suspicion remains that such effects are also present in foreign exchange markets.

[7] One explanation for their finding is that a dealer with an information advantage is unlikely to subsequently transact in the more transparent broker market from where Breedon and Vitale take their data. Instead he will manage his inventory in the opaque direct market to prolong his advantage. The fast, efficient and transparent broker market is more suited to managing inventory positions caused by uninformed order flow.

above is based on data from Citibank, perhaps the most active bank in the foreign exchange market. Bjønnes and Rime (2001) analyse the actions of two dealers in a Scandinavian commercial bank over a trading week, and Mende, Menkhoff and Osler (2004) look at the actions of a small German bank in the euro-dollar market over a four-month period. Both of these papers find customer orders to be important from the dealers' perspectives. Bjønnes and Rime show that their dealers use customer order flow to form their own order placement strategy, but not their own (inter-dealer) quotes.[8] After controlling for dealer inventory, dealers tend to follow the trades of customers (i.e. after a customer buys a currency, dealers tend to buy the currency on the inter-dealer market). Mende, Menkhoff and Osler find that their bank's order flow has predictive power for exchange rates, with a half-life of around fifteen hours. However, these papers do not explicitly differentiate between the alternative explanations for the link between order flows and currency movements. The aim of this paper is to explore further the nature and cause of the relationship between customer order flow and exchange rate changes. Our findings should be considered alongside the complementary ones in Evans and Lyons (2004). Using the Citibank data mentioned above, they show that flows have forecasting power for future macro fundamentals and future spot rates, and that spot rates only slowly impound the information in flows. Evans and Lyons interpret these results as suggesting that flows are part of the process by which low frequency, fundamental information about exchange rates is incorporated into the price.

2. Data description

The data used in this paper come from the Royal Bank of Scotland (RBS). RBS is among the top ten global foreign exchange banks and is probably number one in the pound sterling markets. Customer order flow data are obviously highly confidential and so the data description provided here is necessarily less detailed than usual. However, we hope that readers still get a feel for the nature of the flows across this bank's foreign exchange desks.

---

[8] Yao (1997) reports similar findings from his study of a single US-based dealer over a trading month.

The RBS maintains a 24-hour foreign exchange trading service for its customers. The customer order flow data are aggregated over a 24-hour window from the opening of the Sydney market through the close of the US market (which approximates to midnight to midnight Greenwich Mean Time). The data include all spot transactions entered into by customers against the bank. Thus the data do not include forward deals or deals between the bank and other banks via the inter-dealer markets. The data set begins on 1 August 2002 and ends 29 June 2004, a period of around 460 trading days once (currency-specific) holidays are excluded.

In this paper we use customer order flow data for four currencies: US dollar, euro, Japanese yen and British pound. This implies a set of six bilateral exchange rates and we have order flow figures for each of these. We use this group of currencies because they are among the most heavily traded currencies according to the BIS tri-annual surveys of foreign exchange market activity, and because they are the only set of currencies in our data for which the full set of bilateral exchange rates are traded. We will make use of this mini-system of exchange rates below.

The order flow for each exchange rate is further disaggregated according to the counterparty classification assigned by the bank. There are four categories of customer: non-financial corporates (denoted Corp), unleveraged financials such as mutual funds (Unlev), leveraged financials including hedge funds (Lev) and other financials (Other). The final category is rather heterogeneous but will include the trades of smaller banks that do not have access to the interbank dealer network and trades of central banks. Since central banks do not necessarily trade for profit reasons we differentiate between other financial institutions and profit-maximising financial institutions (leveraged and unleveraged) in our discussions below.

Contemporaneous spot exchange rate data for the corresponding time period were obtained from Norgate Investor Services. The spot exchange rate data include the Sydney opening and New York closing prices, used to calculate daily log changes in exchange rates. Computing daily changes using Sydney open to Sydney open or New York close to New York close did not materially affect any of our results but are available on request.

RBS has asked us not to disclose the magnitude of their customers' gross or net order flows. Instead, Table 1 contains some descriptive statistics of the absolute values of *normalised* net order flows. Absolute total net flows (customer buy orders minus customer sell orders) have been scaled to have a mean of unity for each currency. Net order flows for each counterparty classification are then expressed relative to this, such that the mean absolute corporate net flow in the euro-dollar market is 0.556 times the mean absolute total net flow.

Net order flows are very volatile and in many exchange rate-counterparty classification combinations the standard deviation is greater than the mean absolute net flow. As an illustration, the maximum absolute net order flows from leveraged and unleveraged financials in the pound-yen market were both more than one hundred times the mean flows.

The normalisation (deliberately) masks the relative sizes of the six exchange rate markets, but we can give a ranking based on average daily absolute net order flows in the sample period. The euro-dollar market typically exhibits the largest net order flow, followed by dollar-yen and pound-dollar. Euro-pound, euro-yen and pound-yen cross rates typically see smaller average absolute net order flows.

Table 2 shows that total customer buy and sell order flows are significantly positively autocorrelated for the most frequently traded exchange rates. The other financial institutions (and sometimes non-financial corporate) components of the flows appear to drive this autocorrelation. Leveraged and unleveraged financial institution buy and sell orders are usually less serially correlated. Despite this predictability of gross flows, net order flows are typically not autocorrelated and cumulated net order flows follow random walks.

Table 3 shows that flows from different counterparty classifications typically are not highly correlated. However, flows from other financial institutions are sometimes very negatively correlated with flows from other customer classifications, particularly in the smaller pound-yen market.

3. Regression results

Much of the impetus behind the growth in microstructure research in foreign exchange comes from the simple but controversial correlation between order flow in a given period and the change in the exchange rate over the same period. Finding such a correlation is encouraging because researchers have failed to come up with any other variables that are reliably correlated with short-term exchange rate movements. Finding such a correlation is controversial because it is still not clear whether order flows cause exchange rate changes or vice versa. Researchers sceptical of the microstructural approach to exchange rates worry about the effect of positive feedback trading. Even if the causation is from flows to the exchange rate, it is not straightforward to decide whether this is due to the liquidity or informational effects of order flow. The liquidity effect arises because excess customer demand for a currency would only be supplied by dealers if they were compensated by a shift in the exchange rate. We hope to shed some light on these issues.

3.1 Total order flow and exchange rate changes

As discussed in section 1, there are few customer order flow data sets available for academic research. Ours covers a longer period than most (almost two years) and includes more exchange rate pairs than other data sets. The first step is to establish that there is a correlation between flows and exchange rates in our data set using the following simple regression:

$$\Delta s_t = \beta_0 + \beta_1 x_t + u_t . \qquad \textbf{(1)}$$

The dependent variable is the change in the log of the spot exchange rate and the single independent variable is the total customer order flow (value of customer buys-value of customer sells). A positive β_1 coefficient would suggest that positive order flow into a currency (net buying pressure) is associated with an appreciation of the currency. An intercept term is included but not reported. Its exclusion does not materially affect any of our findings.

Table 4 reports the results of OLS estimation of equation (1) at one-day and one-week horizons for each of our six exchange rates. For the weekly horizon we employ overlapping windows to maximise the amount of information available, and correct the standard errors for the induced serial correlation. At both the daily and weekly frequency we use heteroscedasticity robust standard errors.

Most estimated β_1 coefficients are positive, but only two are significant at the daily frequency and R^2 values are essentially zero. The weekly horizon provides slightly more encouragement, with significant coefficients for three of the six exchange rates and R^2 values as high as sixteen percent. However, the three insignificant coefficients are all negative, and these include the large and liquid euro-dollar and pound-dollar markets. Reversing the direction of the regression (i.e. regressing flows on exchange rate changes) does not alter the significance or sign of any of the coefficients, highlighting the problem of inferring causality from this relationship.

3.2 Disaggregated order flow and exchange rate changes

Estimating regression equation (1) imposes the constraint that the impact of net order flow on the exchange rate is equal for all customer types.[9] This may be reasonable if the correlation between exchange rates and order flow is due to liquidity effects since in this case the nature of the counterparty should be irrelevant – the market maker should adjust his price equally for a trade of a given size from a corporate or financial customer. It may not be a reasonable constraint if the correlation is due to private information since it is conceivable that some types of customers are more informed than others. Carpenter and Wang (2003) discuss the behaviour of customers in foreign exchange markets. They conclude that orders from financial institutions (including central banks) could be expected to contain incremental information, while corporate order flows should not (and they subsequently find evidence supporting these conjectures).

This constraint is relaxed in Table 5 where exchange rate changes are regressed on disaggregated net order flows.

$$\Delta s_t = \beta_0 + \beta_1 x_t^{Corp} + \beta_2 x_t^{Unlev} + \beta_3 x_t^{Lev} + \beta_4 x_t^{Other} + u_t . \qquad \textbf{(2)}$$

The p-values for the LR-test that the coefficients on each component order flow are jointly equal to zero are reported for each regression. These indicate that each regression is significant at both daily and weekly horizons. R^2 values are still relatively low at the daily frequency but are as high as 27 percent over a week.

The influence of different customer categories clearly differs. Non-financial corporate customer flows are significant at the five-percent level in six of the twelve regressions (and in an additional two at the ten-percent level). In each case the

---

[9] Or, for the more sceptical, that counterparty types react equally for a given exchange rate movement.

coefficient is negative (and it is usually negative even when not significant). Profit-maximising financial company flows are always positive. Further, flows from unleveraged financials are significant at the five-percent level in seven regressions, and flows from leveraged financials are significant in five. Coefficients on order flows from other financials are mixed, but are usually positive when significant. It is noticeable that significantly positive coefficients on order flows from other financials (which would include central bank transactions) are always present for yen rates, but not for other currencies. This perhaps reflects the market perception that the Bank of Japan more frequently intervenes in the foreign exchange markets than other central banks.

The coefficients are also economically significant. In the euro-dollar market, for example, a net flow of €1bn into the euro from leveraged financial institutions is associated with a 1.49% rise in the value of the euro over one day, and 1.86% over a week. A similar net flow from non-financial corporates is associated with a fall in the value of the euro of 0.68% over one day (and 0.93% over a week). These numbers are broadly comparable with those found for Citibank's customers. Lyons (2001) reports that a €1bn net flow from leveraged funds [non-financial corporates] is associated with a 0.6% appreciation [0.2% depreciation] of the euro over a month.[10]

Flows in other markets have much higher coefficients. A net flow of €1 billion from leveraged financials in the euro-yen market, for example, would see the euro some four percent higher (although this is only a marginally statistically significant effect). There is some association between the magnitude of the coefficients and the liquidity of each market. Coefficients are relatively small in the very liquid euro-dollar and dollar-yen markets, and are relatively large in the smaller euro-yen and pound-yen cross-rate markets. This could be seen as supporting the liquidity explanation for the correlations.

However, the heterogeneity and broadly systematic pattern of the coefficients suggests that there is some information content in customer order flows. If the relationship between flows and exchange rate changes is due simply to liquidity

[10] Similarly, Froot and Ramadorai's (2004) analysis of State Street Corporation's institutional investor flow data suggest that a €1bn net flow into the euro is associated with an appreciation of the euro of 0.89% over a day, 1.08% over a week and 1% over a month.

effects then there should be no difference between equal sized orders from, for example, a corporate and an unleveraged financial institution. Our results suggest that there is a difference. Specifically, they are consistent with the joint hypothesis that some customer types tend to be more informed than others and that the market as a whole is able to discriminate between flows from different customer types.[11] Our findings parallel those of Ferguson, Mann and Waisburd (2004) who demonstrate that trades from (informed) speculators have much larger price impact than trades from (relatively less informed) hedgers.

We acknowledge that our findings could also be because the nature and degree of feedback trading differs across participants. In particular, perhaps the most appealing explanation of the robustly negative coefficient on corporate customer order flows is that they follow negative feedback rules (i.e. corporates buy the currency that has just depreciated). Dealers and foreign exchange salespeople have suggested to us that corporates often use advantageous short-term exchange rate changes to exchange money for non-speculative reasons (e.g. repatriation of funds). In order to explain the significantly positive coefficients on profit maximising financial institutions' order flows using the feedback explanation, these institutions would have to be following positive feedback trading rules, buying appreciating currencies. This is not totally implausible since many leveraged funds are known to follow momentum-trading strategies.

We address this issue in a simple way by regressing daily disaggegated order flows on lagged exchange rate changes, reporting the results in Table 6.[12] Four out of six coefficients are significantly negative for corporate flows, suggesting that this group of customers responds to prior exchange rate changes in a way consistent with

---

[11] There is one caveat to this assertion. Leveraged funds move large amounts of money and are likely to split their deals between banks. Thus the order flow from leveraged funds observed by our bank is also observed by other banks simultaneously. Deals from other customers are likely to be smaller and/or are unlikely to be split across banks. This could account for the higher coefficient on leveraged flows. However, it cannot reconcile the negative coefficient on corporate flows.

[12] We also attempted to address possible feedback effects using the identification through heteroscedasticity approach of Rigobon and Sack (2003). Unfortunately, this approach could not disentangle the price impacts from the feedback effects, perhaps because the identifying assumption of at least one homoscedastic shock was not supported by the data.

negative feedback. There is also evidence supportive of feedback trading for the other customer classes. However, the evidence is either infrequent (lagged exchange rate changes only seem to affect unleveraged order flows in the dollar-yen market) or mixed (both leveraged and other financials follow positive feedback trading for some currencies but negative for others). These findings do not rule out positive feedback trading by financials within the day, but they are more strongly suggestive of negative feedback trading by corporate customers.

3.3 Cross-market flow effects

Regressions in the form of equation (2) again impose untested restrictions on the nature of the information contained in customer order flows. As specified, they only allow order flow in a particular exchange rate market to affect that exchange rate. The dispersed information model of Evans and Lyons (2002c) suggests that information relevant to the value of a currency may be in the hands of a customer. By trading in a particular exchange rate market this information is revealed and affects other exchange rates. For example, a customer may have value-relevant information regarding the euro. By trading euro-dollar, this information is partly revealed, directly to the euro-dollar market and indirectly to the 'related' euro-pound and euro-yen markets.[13] To the extent that this same information is value-relevant to non-euro currencies we might also expect seemingly 'unrelated' exchange rates such as pound-yen to react to the order flow.

In this framework, a single exchange rate could be expected to react not only to (disaggregated) flows in its own market, but also to flows in all other exchange rate markets whether related or not. This leads to the regression equation (3), shown here with the change in the euro-dollar exchange rate as dependent variable:

[13] The pound-dollar and dollar-yen markets are also related to the euro-dollar through the US dollar side of the bargain.

$$
\begin{aligned}
\Delta s_{\$/€t} &= \beta_0 + \beta_{1\$/€} x^{Corp}_{\$/€t} + \beta_{2\$/€} x^{Unlev}_{\$/€t} + \beta_{3\$/€} x^{Lev}_{\$/€t} + \beta_{4\$/€} x^{Other}_{\$/€t} \\
&\quad + \sum_R \left(\beta_{1R} x^{Corp}_{Rt} + \beta_{2R} x^{Unlev}_{Rt} + \beta_{3R} x^{Lev}_{Rt} + \beta_{4R} x^{Other}_{Rt} \right) \\
&\quad + \beta_{1£/Y} x^{Corp}_{£/Yt} + \beta_{2£/Y} x^{Unlev}_{£/Yt} + \beta_{3£/Y} x^{Lev}_{£/Yt} + \beta_{4£/Y} x^{Other}_{£/Yt} + u_t \\
R &= \{\$/Y, \$/£, €/Y, €/£\}
\end{aligned} \tag{3}
$$

The first line of equation (3) allows changes in the euro-dollar rate to be related to flows in the euro-dollar market (such that this component of the regression is equivalent to equation (2)). The second line allows the euro-dollar rate to be related to flows in the four related markets (other dollar bilateral rates and other euro bilateral rates) while the third line allows flows in the unrelated pound-yen market to matter.

Tests of this less restrictive model of the importance of order flow are reported in Table 7. At the daily frequency, all six exchange rates react to 'own' and 'related' order flows. The euro-dollar rate even reacts to flows in the 'unrelated' pound-yen market. At a weekly frequency, five out of six exchange rates are influenced by their own order flows at the five-percent significance level (the remaining euro-pound rate is significant at the 8% level), and all six are influenced by related order flows. Five are even influenced by order flows in the unrelated exchange rate – euro-dollar (again), pound-yen and dollar-yen react to unrelated flows at high significance levels, and euro-pound and pound-dollar at more marginal ones.

While not conclusive, this again suggests that liquidity effects are not the main cause of the correlation between flows and exchange rate changes. Dealers have told us that large, experienced teams of dealers do not typically manage exposures on a portfolio basis but instead do so exchange rate by exchange rate.[14] The management of the flows faced by the euro-pound dealer does not typically influence the risk-management actions of the euro-dollar trader. However, communication across the dealing desk is such that information about flows in other exchange rates is exchanged and, according to Table 6, is related to exchange rate movements.

---

[14] Smaller, less experienced teams may follow a portfolio approach in the presence of a "hands on" head trader.

4. Probability of informed trading

Results so far suggest that the high contemporaneous correlation between order flows and exchange rate changes may be due to information asymmetries. However, we acknowledge that since the simple regressions are reduced form, firm conclusions are impossible and that, in particular, feedback trading could still be at the root of the correlations. In this section we look more closely at the nature of the order flows using a method now widely accepted in the equity market microstructure literature, which purports to determine the probability of information-based trading (PIN). While very simple, this measure has been shown to explain a number of information-based regularities in equity markets.[15] For example, Easley, Kiefer, O'Hara and Paperman (1996) show that low-volume stocks face higher probabilities of informed trading and that this can explain the higher spreads charged on such stocks. More recently, Easley, Hvidkjaer and O'Hara (2002) show that equities with greater private information command a risk premium. A ten-percent increase in PIN is associated with an increase in annual expected returns of 2.5%. To our knowledge, ours is the first paper to apply the PIN model to exchange rates.

4.1 The probability of information-based trading model

The PIN model was developed by Easley, Kiefer and O'Hara (1996, 1997a, 1997b). They demonstrate how a simple structural model can provide specific estimates of the risk of information-based trading in an asset. The model is based on the trading game played by a market-maker and customers, repeated over independent and identically distributed trading intervals $i = 1,\dots, I$. At the start of each trading interval nature decides whether there is new information available. New information is available with probability α. This new information is a signal regarding the underlying asset value, and can be good news for the asset, suggesting a high price, or bad news, suggesting a low price. Conditional on new information occurring, good news

---

[15] It is not universally accepted that the PIN model is capturing customer informedness as intended. Aktas, de Bodt, Declerck and Van Oppens (2003) find counter-intuitive results when using PIN around merger announcements, although they apply the model in a limit-order book environment (the Paris bourse) rather than the market-maker setting that the original theory assumes.

happens with probability (1-δ) and bad news with probability δ. Customers arrive according to Poisson processes throughout the trading interval and the market maker sets buy and sell prices at each point in time and executes orders as they arrive. Some customers are able to observe the new information, and are termed informed. Informed customers arrive at a rate μ (in information periods) and buy if they have observed good news and sell if they have observed bad news. Other customers and, crucially, the market maker are not able to observe the new information. Uninformed customers arrive and buy at rate ε_b and arrive and sell at rate ε_s. For simplicity we assume these two rates are equal to ε. If an order arrives at time *t*, the market maker observes the trade and uses this information to update his beliefs about the underlying value of the asset, setting new prices accordingly.

Gross and net order flows allow the econometrician to estimate the key parameters of this model. The total trades made per interval (*TT* = buys plus sells) equals the sum of the Poisson arrival rates of informed and uninformed customers.

$$TT = \underbrace{\alpha(1-\delta)(\varepsilon+\mu+\varepsilon)}_{\text{good news interval}} + \underbrace{\alpha\delta(\mu+\varepsilon+\varepsilon)}_{\text{bad news interval}} + \underbrace{(1-\alpha)(\varepsilon+\varepsilon)}_{\text{no news interval}} = \alpha\mu + 2\varepsilon$$

The trade imbalance (*K* = sells – buys) is such that

$$K = \alpha\mu(2\delta - 1)$$

More informatively, the absolute value of the net order flow, $|K|$ approximates to $\alpha\mu$ for large enough levels of μ.

Easley, Kiefer and O'Hara show that in trading interval *j*, conditional on the parameter vector $\Theta = [\alpha, \delta, \mu, \varepsilon]^{\mathrm{T}}$, the probability of observing *B* buys and *S* sells is given by

$$\Pr[y_j = (B,S)|\Theta] = \alpha(1-\delta)e^{-(\mu+2\varepsilon)}\frac{(\mu+\varepsilon)^B(\varepsilon)^S}{B!S!} + \alpha\delta e^{-(\mu+2\varepsilon)}\frac{(\mu+\varepsilon)^S(\varepsilon)^B}{B!S!} + (1-\alpha)e^{-2\varepsilon}\frac{(\varepsilon)^{(B+S)}}{B!S!}$$

Because of the assumption of identically-distributed and independent trading intervals, the likelihood function is the product of this probability density over trading intervals.

The model allows the econometrician to apply maximum likelihood techniques to observed count data on the number of buys and sells to make inference about the division of trade between informed and uninformed customers. Because of the daily frequency of our data we are forced to assume the trading interval corresponds to a single day. While we recognise that this is less than ideal, we note that this approach has proved successful using daily data on equities where it seems equally implausible that there is at most a single information event each day. The model is usually estimated using the number rather than the value of buy and sell transactions. Unfortunately, neither the number of transactions nor the values of individual transactions are available to us. We make the simple assumption that each transaction is for one million units (euros, pounds or dollars depending on the exchange rate being considered) and transform our daily value series into daily counts of transactions.

4.2 PIN estimation results

The results of estimating the PIN model for each exchange rate and customer type are presented in Table 8. Since the estimates of μ and ε would reveal the gross trading volume of the bank providing the flow data we do not report the specific coefficient estimates. Instead, where convergence of the maximum likelihood algorithm is possible, we report the probability of informed trading, PIN:

$$PIN = \frac{\alpha\mu}{\alpha\mu + 2\varepsilon}$$

These values are quite precisely estimated. The *t*-statistics on estimates of μ and ε are usually greater than 100 and never below 10, while *t*-statistics on α are always in excess of 5.[16,17] PINs range from 0.09 to 0.63 with a mean of 0.31. This compares

[16] Easley, Hvidkjaer and O'Hara (2002) note that information on μ and ε accumulates at a rate approximately equal to the number of trade outcomes while information on α (and δ) accumulates at a rate equal to the square root of the number of trading days, which explains the difference in precision.

[17] These *t*-statistics should not be taken at face value. Hasbrouck (2004) demonstrates the difficulty in obtaining precise estimates of parameter values from mixtures of distributions such as the PIN model, but argues that the PIN itself is likely to be well-identified.

with a mean PIN of 0.21 for the sample of all ordinary stocks listed on the NYSE used in Easley, Hvidkjaer and O'Hara (2002).

The first value in Table 8 suggests that if the first trade in the daily trading interval is by a corporate customer, the bank's euro-dollar market maker should expect the probability that the customer is informed to be 21%. If that first trade is by an unleveraged financial institution then the market maker should assume a probability in excess of 36%.

Significantly, a clear pattern emerges from the PINs estimated for customer groups. For every exchange rate, the PIN estimated for corporate customers is less than that for leveraged and unleveraged financial institutions.[18] In four out of six cases, the probability of informed trading by leveraged financial institutions is greater than that of the unleveraged financial institutions.

The PIN is determined by the value of $\alpha\mu$ relative to 2ε, that is the arrival rate of informed trades relative to uninformed trades. Other things equal, PIN will be high for a customer class if information events happen frequently (α is high), informed traders arrive frequently (μ is high), or uninformed traders arrive infrequently (ε is low). Which of these terms lead the PINs for financial institutions to be so much higher than the PINs of corporates?

In most cases we find that the estimate of α is higher for financials than for corporates.[19] In the euro-pound market, for example, financial customers observe an information event with probability 0.2, compared to 0.04 for corporates. This suggests that private information is playing a part. Financial customers appear to be looking at a wider information set (or looking closer at the same information set) and as a result observe price-relevant signals more often.

---

[18] PINs for other financial institutions are typically relatively small. In two cases convergence was impossible because of occasional extremely high numbers of trades compared to normal levels. Easley, Hvidkjaer and O'Hara (2004) discuss in more detail some problems encountered when estimating PINs.

[19] We find higher α estimates for financial customers than corporates for each exchange rate with the exception of pound-yen.

The ratio of μ to 2ε gives the arrival rate of informed customers relative to uninformed customers conditional on an information event having occurred. This ratio is higher for financial customers than corporate customers for every exchange rate except euro-yen. So, for example, even if informed corporate clients in the euro-dollar market have observed an information event (which happens 26.3% of the time), their trades are not much in excess of the trades of uninformed corporates since the ratio of μ to 2ε is 1.02. The ratio for unleveraged institutions is in excess of 1.76 (and informed leveraged customers observe signals almost 33% of the time).

It appears, therefore, that both the higher relative arrival rate of informed to uninformed customers and the greater probability that they have observed a signal contribute to the higher PIN for financial customers.[20]

5. Price-impact and PIN results consistency

Section 3.2 suggests that the price impacts of net order flows from different customer classifications differ in a broadly systematic way. Corporate customers' net order flows have zero or negative correlation with exchange rates. Leveraged and unleveraged financial customers' net order flows have significantly positive price impact (in using this term we assume causation runs from flows to rates). We argue that this pattern reflects the different private information contained in order flows from different customer groups. These different price impacts across customer classifications are not supportive of the idea that the price impact of net order flow is due to liquidity effects. Section 3.3 suggests that order flows in one exchange rate market can affect the spot rate in another market (although the two markets will typically share a currency). This is again not supportive of the liquidity effect explanation. However, we still cannot rule out feedback trading as the root cause of the correlation between flows and exchange rate.

---

[20] Hasbrouck (2004) shows that the PIN estimate is driven by the product of α and μ and that it may be difficult to separate these two terms empirically – the same value of PIN can be generated by a low α and high μ, or a high α and low μ. However, the variation across categories suggests that both the α and μ terms are higher for financial customers than for corporates.

Section 4 suggests that different customer classifications have different probabilities of being informed. These estimates are based solely on order flow data and do not utilise the relevant exchange rates. As a final test of the hypothesis that order flow impacts the exchange rate because of private information, we now examine the consistency of the results of sections 3 and 4.

The correlation between the PIN of the 22 exchange rate-customer classification combinations given in Table 7 and the coefficients from the daily regressions reported in Table 5 is 0.69. Using the coefficients from the weekly regressions raises this marginally to 0.71.[21] There is a strongly positive relationship between the estimated probability of order flows being information-related, and the nature of the link between these flows and the exchange rate.[22] This is our final piece of evidence that customer order flows appear to contain information relevant to exchange rate pricing.

6. Conclusions

There is little evidence that standard macroeconomic model of exchange rates have anything to say about high frequency movements beyond the impact of news announcements. There is a growing literature that suggests the microstructure approach has something more positive to contribute. Foremost in this literature is the strong contemporaneous correlation between order flows and exchange rate changes. Alternative explanations for this exist. The three most common are that participants in the market may follow high frequency feedback trading rules, that risk-averse dealers may need to be compensated for holding unwanted inventory (the liquidity effect), or that order flows may contain price-relevant information. This paper uses a new database of daily customer order flows from a large commercial bank to differentiate between these explanations.

---

[21] Recognising that some of the exchange rates analysed are more liquid than others (and so the link between flows and exchange rates is likely to differ across markets) we also estimate the correlations for each exchange rate. Every correlation is positive, ranging from 0.37 (dollar-yen, weekly) to 0.99 (pound-yen, weekly). The average correlation is 0.72.

[22] This contrasts markedly with the small but significant correlation of 0.0338 between price impact and PIN calculated for a selection of equities on US stock exchanges by Dennis and Weston (2001).

One of the main advantages of this database is that it classifies order flows according to the nature of the customer. This allows us to confirm the findings of Lyons (2001) that order flows from different customer types have different correlations with exchange rate changes. Flows from profit-seeking financial institutions are positively correlated, while flows from non-financial corporates are typically negatively correlated. The trading motivation for the final customer category – "other financial institutions" – is not clear and results for this group are mixed. A second advantage of the data is that they cover several exchange rate pairs. We examine a mini-system of six rates and find that flows in one exchange rate pair are usually associated with changes in other exchange rates. Together, these two pieces of evidence lead us to rule out the liquidity explanation. Foreign exchange dealers do not typically manage inventories on a portfolio basis (so flows in another exchange rate market should not matter) and if they demand a risk premium for unwanted inventory they should demand a similar risk premium no matter where the inventory comes from.

A final advantage of the data is that we have gross flows (the value of buy orders and sell orders) rather than just net flows (buys minus sells). This allows us to estimate Easley *et al*'s PIN measure of the informedness of each customer classification, a technique widely used in the equity market literature. These estimates suggest that order flows from non-financial corporates are usually much less likely to contain information relevant to the value of a currency than leveraged or unleveraged financial institutions. Two factors lie behind this. First, financial institutions appear to observe price-relevant signals more frequently than corporates, a plausible result since financial institutions are supposed to be actively monitoring markets seeking profit opportunities. Second, conditional on receiving a price-relevant signal, financial institutions trade more aggressively than corporates, again quite plausible since financial institutions have capital allocated for speculative trading while corporates typically do not.

Finally, we show that the nature of the order flow-exchange rate change correlation is itself very positively correlated with the estimated probability of the order flow being from an informed source. We conclude from this that order flow from profit maximising financial institutions – leveraged and unleveraged financials – does indeed contain price relevant information and that this explains the positive correlation between their order flows and exchange rate changes. We are left with the puzzling negative correlation between non-financial corporate order flows and exchange rates. We think that the most plausible explanation for this is that corporates follow negative feedback trading strategies and provide some evidence that corporate flows respond to lagged exchange rate movements.

References

Aktas, Nihat, Eric de Bodt, Fany Decleck and Herve Van Oppens (2003) “PIN? Some evidences around corporate events,” *mimeo,* Universite Catholique de Louvain.

Bjønnes, Geir H., and Dagfinn Rime (2001) “Customer trading and information in foreign exchange markets,” *mimeo*, Norwegian School of Management.

Bjønnes, Geir H., and Dagfinn Rime (2003) “Dealer behavior and trading systems in foreign exchange markets,” *Journal of Financial Economics*, forthcoming.

Breedon, Francis and Paulo Vitale (2004) “An empirical study of liquidity and information effects of order flow on exchange rates,” *mimeo*, Imperial College London.

Carpenter, Andrew and Jianxin Wang (2003) “Sources of private information in FX trading,” *mimeo,* University of New South Wales.

Cheung, Yin-Wong, Menzie D. Chinn and Ian W. Marsh (2000) “How do UK-based foreign exchange dealers think their market operates?” *NBER working paper 7524*.

Cohen, Benjamin H., and Hyun Song Shin (2003) “Positive feedback trading under stress: Evidence from the US Treasury securities market,” *mimeo,* London School of Economics.

Danielsson, Jon and Ryan Love (2004) “Feedback trading,” *mimeo*, London School of Economics.

Danielsson, Jon, Richard Payne and Jinhui Luo (2002) “Exchange rate determination and inter-market order flow effects,” *mimeo,* London School of Economics.

Dennis, Patrick J., and James P. Weston (2001) “Who’s informed? An analysis of stock ownership and informed trading,” *mimeo,* University of Virginia.

Easley, David, Nicholas M. Kiefer, and Maureen O’Hara (1996) “Cream-skimming or profit sharing? The curious role of purchased order flow,” *Journal of Finance*, 51, 811-833.

Easley, David, Nicholas M. Kiefer, and Maureen O’Hara (1997a) “One day in the life of a very common stock,” *Review of Financial Studies*, 10, 805-835.

Easley, David, Nicholas M. Kiefer, and Maureen O’Hara (1997b) “The information content of the trading process,” *Journal of Empirical Finance*, 4, 159-186.

Easley, David, Nicholas M. Kiefer, Maureen O’Hara, and Joseph Paperman (1996) “Liquidity, information, and less-frequently traded stocks,” *Journal of Finance*, 51, 1405-1436.

Easley, David, Soeren Hvidkjaer, and Maureen O’Hara (2002) “Is information risk a determinant of asset returns?” *Journal of Finance*, 57, 2185-2221.

Easley, David, Soeren Hvidkjaer, and Maureen O’Hara (2004) “Factoring information into returns,” *mimeo*, University of Maryland.

Evans, Martin D., and Richard K. Lyons (2002a) “Order flow and exchange rate dynamics,” *Journal of Political Economy*, 102, 170-180.

Evans, Martin D., and Richard K. Lyons (2002b) “Time-varying liquidity in foreign exchange,” *Journal of Monetary Economics*, 49, 1025-1051.

Evans, Martin D., and Richard K. Lyons (2002c) “Financial integration and FX trading,” *Journal of International Money and Finance*, 21, 807-831.

Evans, Martin D., and Richard K. Lyons (2004) “Exchange rate fundamentals and order flow,” *mimeo*, Georgetown University.

Fan, Mintao and Richard K. Lyons (2003) “Customer trades and extreme events in foreign exchange,” in *Essays in Honor of Charles Goodhart*, Paul Mizen (ed.), Edward Elgar: Northampton, MA, USA, pages 160-179.

Ferguson, Michael F., Steven C. Mann, and Andrew C. Waisburd (2004) “Discriminating between investor types on the exchange floor,” *mimeo*, University of Cincinnati.

Froot, Kenneth A.,and Tarun Ramadorai (2004) “Currency returns, intrinsic value, and institutional investor flows,” *Journal of Finance*, forthcoming.

Hasbrouck, Joel (2004) “Empirical market microstructure,” *mimeo*, New York University.

Lyons, Richard K. (1995) “Tests of microstructural hypotheses in the foreign exchange market,” *Journal of Financial Economics*, 39, 321-351.

Lyons, Richard K. (1997) “A simultaneous trade model of the foreign exchange hot potato” *Journal of International Economics*, 42, 275-298.

Lyons, Richard K. (1998) “Profits and position control: A week of FX dealing,” *Journal of International Money and Finance*, 17, 97-115.

Lyons, Richard K. (2001) *The microstructure approach to exchange rates*, MIT Press.

MacDonald, Ronald and Ian W. Marsh (2004) “Currency spillovers and tri-polarity: A simultaneous model of the US dollar, German mark and Japanese yen,” *Journal of International Money and Finance*, 23, 99-111.

Mark, Nelson (1985) “Exchange rates and fundamentals: Evidence on long-horizon predictability,” *American Economic Review*, 85, 201-218.

Mende, Alexander, Lukas Menkhoff and Carol L. Osler (2004) “Asymmetric information and the cross-section of currency spreads,” *mimeo*, University of Hannover.

Payne, Richard (2003) “Informed trade in spot foreign exchange markets: An empirical investigation,” *Journal of International Economics*, 61, 307-329.

Rigobon, Roberto, and Brian Sack (2003) “Measuring the reaction of monetary policy to the stock market,” *Quarterly Journal of Economics*, 118, 639-669.

Yao, Jian (1997) “Market making in the interbank foreign exchange market,” *mimeo*, New York University.

Table 1

This table presents descriptive statistics for absolute values of net order flows (buys minus sells) for each bilateral currency pair. Details are given for total flows and for the four categories of customer described in the text. To mask the value of RBS's order flows we normalise the mean total order flow to equal unity for each currency pair. Thus, the mean corporate customer absolute net order flow is 0.556 times the total absolute net order flow for the euro-dollar market.

	Euro-dollar					**Euro-yen**				
	Total	Corp	Unlev	Lev	Other	Total	Corp	Unlev	Lev	Other
Mean	1.000	0.556	0.297	0.318	0.755	1.000	0.401	0.226	0.170	0.765
Median	0.716	0.405	0.158	0.211	0.550	0.685	0.173	0.109	0.076	0.509
Minimum	0.004	0.001	0.000	0.000	0.002	0.005	0.000	0.000	0.000	0.002
Maximum	6.716	3.712	4.782	2.756	4.978	10.416	9.216	2.635	2.030	6.291
Std Deviation	0.937	0.553	0.479	0.353	0.750	1.099	0.813	0.338	0.248	0.814
	Euro-pound					**Pound-yen**				
	Total	Corp	Unlev	Lev	Other	Total	Corp	Unlev	Lev	Other
Mean	1.000	0.619	0.225	0.275	0.485	1.000	0.406	0.341	0.139	0.731
Median	0.687	0.425	0.100	0.086	0.335	0.554	0.186	0.079	0.000	0.364
Minimum	0.001	0.004	0.000	0.000	0.002	0.000	0.001	0.000	0.000	0.000
Maximum	8.532	7.849	3.409	8.340	5.544	36.713	21.598	35.614	14.083	36.409
Std Deviation	1.080	0.761	0.399	0.656	0.536	2.211	1.122	1.760	0.742	2.508
	Pound-dollar					**Dollar-yen**				
	Total	Corp	Unlev	Lev	Other	Total	Corp	Unlev	Lev	Other
Mean	1.000	0.516	0.318	0.319	0.622	1.000	0.449	0.329	0.233	0.734
Median	0.713	0.341	0.172	0.167	0.429	0.765	0.289	0.208	0.141	0.500
Minimum	0.009	0.001	0.001	0.000	0.004	0.005	0.002	0.000	0.000	0.002
Maximum	8.604	4.140	4.469	3.461	4.640	7.438	5.831	3.117	2.116	7.676
Std Deviation	0.969	0.600	0.481	0.420	0.633	0.965	0.543	0.401	0.277	0.818

Table 2

This table provides autocorrelation coefficients for buy orders, sell orders, net orders (buy minus sells), and absolute net orders for each bilateral currency pair. Coefficients are given for the total of all customers and for the four categories of customer described in the text. ADF denotes an augmented Dickey-Fuller test for non-stationarity in the cumulated net order flow series.

	Euro-dollar					**Euro-yen**				
	Total	Corp	Unlev	Lev	Other	Total	Corp	Unlev	Lev	Other
Buys	0.510	0.098	0.257	0.029	0.611	0.177	0.050	0.062	0.000	0.228
Sells	0.516	0.074	0.163	0.048	0.635	0.153	0.017	0.107	-0.008	0.193
Net	0.029	-0.001	0.054	-0.123	-0.014	0.116	0.000	0.028	-0.027	0.159
Abs(Net)	0.011	0.055	0.047	0.087	0.077	0.040	0.015	0.004	-0.017	0.129
ADF	-2.156	0.436	-1.856	-2.482	-2.140	-0.418	-1.001	-1.985	1.045	-1.292

	Euro-pound					**Pound-yen**				
	Total	Corp	Unlev	Lev	Other	Total	Corp	Unlev	Lev	Other
Buys	0.113	0.102	0.042	-0.006	0.105	0.007	-0.006	0.021	0.156	0.017
Sells	0.043	0.015	0.051	-0.006	0.078	0.023	0.042	-0.014	0.054	0.069
Net	0.135	0.098	-0.022	0.012	-0.008	0.047	-0.014	-0.004	0.063	0.036
Abs(Net)	0.051	0.011	0.012	0.003	-0.069	0.049	-0.015	-0.009	0.039	0.017
ADF	2.657	3.215	-0.836	-1.365	-0.273	-0.984	-0.683	-0.660	-1.793	-2.003

	Pound-dollar					**Dollar-yen**				
	Total	Corp	Unlev	Lev	Other	Total	Corp	Unlev	Lev	Other
Buys	0.166	0.142	0.024	0.059	0.286	0.399	0.110	0.085	0.008	0.432
Sells	0.217	0.157	0.005	0.045	0.279	0.417	0.096	0.152	0.009	0.462
Net	-0.032	0.032	-0.031	-0.061	0.015	0.144	0.081	0.063	0.051	0.070
Abs(Net)	0.083	-0.077	0.014	0.065	0.090	0.155	0.120	-0.032	0.049	0.145
ADF	-0.861	-1.194	-1.121	-1.441	-1.159	-1.687	-1.226	-2.024	-1.170	-2.074

Table 3

This table provides cross-correlation coefficients for each bilateral currency pair and for the four different categories of customer described in the text.

	Euro-dollar			Euro-yen		
	Corp	Unlev	Lev	Corp	Unlev	Lev
Unlev	0.042			-0.001		
Lev	0.021	0.052		0.019	0.070	
Other	-0.152	-0.052	-0.173	-0.068	-0.040	0.019
	Euro-pound			**Pound-yen**		
	Corp	Unlev	Lev	Corp	Unlev	Lev
Unlev	0.085			0.001		
Lev	-0.009	0.027		0.024	0.005	
Other	-0.018	0.047	-0.156	-0.013	-0.557	-0.233
	Pound-dollar			**Dollar-yen**		
	Corp	Unlev	Lev	Corp	Unlev	Lev
Unlev	0.038			0.054		
Lev	0.044	0.064		-0.064	0.143	
Other	-0.062	-0.010	-0.109	-0.192	0.024	0.045

Table 4

This table reports the results of OLS regressions of the form: $\Delta s_t = \beta_0 + \beta_1 x_t + u_t$. The dependent variable is the log change in the relevant spot exchange rate over the relevant interval, and the explanatory variable is the total net customer order flow in that exchange rate market during the same interval. The top half of the table reports the one-day interval results. The bottom half uses overlapping five-day intervals with standard errors corrected for the induced serial correlation in the residual. The standard errors for both regressions are robust to heteroscedasticity. Bold p-values denote coefficients significant at the five-percent level.

Daily	**Coefficient ($\times 10^{10}$)**	**Std Error ($\times 10^{10}$)**	**t-statistic**	**p-value**	R^2
Euro-Dollar	-0.0353	0.1509	0.234	0.815	0.000
Euro-Yen	0.9732	0.3768	2.583	**0.009**	0.014
Euro-Pound	0.0435	0.1818	0.239	0.811	0.000
Pound-Yen	0.1692	0.1348	1.254	0.209	0.008
Pound-Dollar	0.2958	0.2611	1.133	0.257	0.003
Dollar-Yen	0.7257	0.2117	3.428	**0.001**	0.043
Weekly					
Euro-Dollar	-0.1979	0.2053	0.964	0.335	0.005
Euro-Yen	2.3775	0.6493	3.662	**0.000**	0.100
Euro-Pound	-0.1287	0.2761	0.466	0.641	0.001
Pound-Yen	2.6942	1.2411	2.171	**0.029**	0.024
Pound-Dollar	-0.6212	0.5258	1.181	0.237	0.013
Dollar-Yen	1.2482	0.2582	4.834	**0.000**	0.161

Table 5

This table reports the results of OLS regressions of the form:

$$\Delta s_t = \beta_0 + \beta_1 x_t^{Corp} + \beta_2 x_t^{Unlev} + \beta_3 x_t^{Lev} + \beta_4 x_t^{Other} + u_t.$$

The dependent variable is the log change in the relevant spot exchange rate over the relevant interval. The explanatory variables are net order flows in that exchange rate market during the same interval disaggregated by customer category as noted in the text. The LR test restricts $\beta_1=\beta_2=\beta_3=\beta_4=0$. Bold p-values denote coefficients or test statistics significant at the five-percent level.

	Daily				**Weekly**			
Euro-Dollar	**Corp**	**Unlev**	**Lev**	**Other**	**Corp**	**Unlev**	**Lev**	**Other**
Coefficient ($\times 10^{10}$)	-0.6859	0.7062	1.4944	-0.1992	-0.9355	0.4694	1.8584	-0.5353
Std Error ($\times 10^{10}$)	0.2604	0.3684	0.3567	0.1785	0.4407	0.4852	0.6312	0.2207
t-statistic	2.635	1.917	4.189	1.116	2.123	0.967	2.944	2.423
p-value	**0.008**	0.055	**0.000**	0.265	**0.034**	0.333	**0.003**	**0.015**
R^2	0.064				0.097			
LR test p-value	**0.000**				**0.000**			
Euro-Yen								
Coefficient ($\times 10^{10}$)	-0.3736	3.6232	4.0213	1.2132	0.7506	4.6024	5.8178	2.5066
Std Error ($\times 10^{10}$)	0.6031	1.1670	2.3866	0.5326	1.2862	3.6951	4.1431	0.8061
t-statistic	0.619	3.105	1.685	2.278	0.584	1.246	1.404	3.109
p-value	0.536	**0.002**	0.092	**0.023**	0.559	0.213	0.160	**0.002**
R^2	0.041				0.124			
LR test p-value	**0.000**				**0.001**			
Euro-Pound								
Coefficient ($\times 10^{10}$)	-0.5776	0.7759	0.9308	-0.0906	-0.9745	0.4133	1.1436	0.2270
Std Error ($\times 10^{10}$)	0.3084	0.5516	0.2902	0.3568	0.4026	1.1803	0.4523	0.5268
t-statistic	1.873	1.407	3.208	0.025	2.420	0.350	2.529	0.431
p-value	0.061	0.159	**0.001**	0.979	**0.016**	0.726	**0.011**	0.666
R^2	0.027				0.055			
LR test p-value	**0.002**				**0.011**			
Pound-Yen								
Coefficient ($\times 10^{10}$)	-2.3533	2.6446	4.6987	2.9553	-2.6062	4.6908	14.160	3.3830
Std Error ($\times 10^{10}$)	1.1433	1.1058	4.0365	0.8206	2.1212	1.4636	7.6556	1.6274
t-statistic	2.058	2.392	1.164	3.601	1.229	3.205	1.850	2.079
p-value	**0.039**	**0.016**	0.244	**0.000**	0.219	**0.001**	0.064	**0.038**
R^2	0.025				0.090			
LR test p-value	**0.000**				**0.001**			
Pound-Dollar								
Coefficient ($\times 10^{10}$)	-1.0754	1.9879	2.6536	-0.2581	-2.3521	3.3405	0.9956	-1.3077
Std Error ($\times 10^{10}$)	0.3675	0.5322	0.7366	0.4420	0.6340	1.0268	1.4374	0.6754
t-statistic	2.926	3.735	3.603	0.584	3.709	3.253	0.693	1.936
p-value	**0.003**	**0.000**	**0.000**	0.559	**0.000**	**0.001**	0.488	0.053
R^2	0.078				0.189			
LR test p-value	**0.000**				**0.000**			

Table 5 – *continued*

Dollar-Yen

Coefficient ($\times 10^{10}$)	-0.6761	1.6461	0.8346	0.8832	-0.6431	3.1379	0.4065	1.1547
Std Error ($\times 10^{10}$)	0.3797	0.4802	0.6338	0.2603	0.4386	0.6117	0.8589	0.2896
t-statistic	1.781	3.428	1.317	3.393	1.466	5.129	0.473	3.987
p-value	0.075	**0.001**	0.188	**0.001**	0.143	**0.000**	0.636	**0.000**
R^2	0.095				0.274			
LR test p-value	**0.000**				**0.000**			

Table 6

This table reports the results of OLS regressions of the form: $x_t = \beta_0 + \beta_1 \Delta s_{t-1} + u_t$. The dependent variable is the net customer order flow in an exchange rate market disaggregated by customer category, and the explanatory variable is the log change in the relevant spot exchange rate during the previous trading day. Bold p-values denote coefficients significant at the five-percent level.

Euro-Dollar	**Corp**	**Unlev**	**Lev**	**Other**
Coefficient ($\times 10^{10}$)	0.00060	0.00122	-0.00960	-0.01712
Std Error ($\times 10^{10}$)	0.00920	0.00752	0.00484	0.01393
t-statistic	0.0655	0.1624	-1.9862	-1.2292
p-value	0.9478	0.8710	**0.0470**	0.2190
Euro-Yen				
Coefficient ($\times 10^{10}$)	-0.00637	0.00091	0.00261	0.02141
Std Error ($\times 10^{10}$)	0.00309	0.00219	0.00107	0.00465
t-statistic	-2.0577	0.4164	2.4480	4.6006
p-value	**0.0396**	0.6771	**0.0144**	**0.0000**
Euro-Pound				
Coefficient ($\times 10^{10}$)	-0.02976	-0.00029	0.00333	0.00527
Std Error ($\times 10^{10}$)	0.00744	0.00395	0.00492	0.00884
t-statistic	-4.0023	-0.0723	0.6776	0.5955
p-value	**0.0001**	0.9424	0.4980	0.5515
Pound-Yen				
Coefficient ($\times 10^{10}$)	0.00008	0.00080	0.00106	0.00410
Std Error ($\times 10^{10}$)	0.00091	0.00060	0.00044	0.00128
t-statistic	0.0930	1.3345	2.4012	3.1989
p-value	0.9259	0.1820	**0.0163**	**0.0014**
Pound-Dollar				
Coefficient ($\times 10^{10}$)	-0.01497	-0.00119	-0.00286	-0.02115
Std Error ($\times 10^{10}$)	0.00515	0.00422	0.00366	0.01016
t-statistic	-2.9094	-0.2824	-0.7795	-2.0815
p-value	**0.0036**	0.7776	0.4357	**0.0374**
Dollar-Yen				
Coefficient ($\times 10^{10}$)	-0.01332	0.02171	0.00575	0.02173
Std Error ($\times 10^{10}$)	0.00719	0.00525	0.00378	0.01113
t-statistic	-1.8517	4.1360	1.5231	1.9533
p-value	0.0641	**0.0000**	0.1277	0.0508

Table 7

This table reports p-values associated with exclusion restrictions on OLS regressions of the form:

$$\Delta s_{\$/€t} = \beta_0 + \beta_{1\$/€} x_{\$/€t}^{Corp} + \beta_{2\$/€} x_{\$/€t}^{Unlev} + \beta_{3\$/€} x_{\$/€t}^{Lev} + \beta_{4\$/€} x_{\$/€t}^{Other}$$
$$+ \sum_R \left(\beta_{1R} x_{Rt}^{Corp} + \beta_{2R} x_{Rt}^{Unlev} + \beta_{3R} x_{Rt}^{Lev} + \beta_{4R} x_{Rt}^{Other} \right)$$
$$+ \beta_{1£/Y} x_{£/Yt}^{Corp} + \beta_{2£/Y} x_{£/Yt}^{Unlev} + \beta_{3£/Y} x_{£/Yt}^{Lev} + \beta_{4£/Y} x_{£/Yt}^{Other} + u_t$$
$$R = \{\$/Y, \$/£, €/Y, €/£\}$$

The dependent variable is the log change in the relevant spot exchange rate (in this example, the euro-dollar rate) over the relevant interval. The explanatory variables are net order flows in all exchange rate markets during the same interval disaggregated by customer category. The column headed “Own” restricts the beta coefficients to be zero in the first line of the equation (excluding the intercept). The column headed “Related” (“Unrelated”) restricts the beta coefficients to be zero in the second (third) row of the equation. The final four columns restrict the coefficients on order flows of each of the four categories of customer to be zero in all currency markets. The top half of the table reports the one-day interval results. The bottom half uses overlapping five-day intervals with standard errors corrected for the induced serial correlation in the residual. The standard errors for both regressions are robust to heteroscedasticity. Bold p-values denote coefficients or test statistics significant at the five-percent level.

		LR-test p-values						
Daily	R^2	**Own**	**Related**	**Unrelated**	**Corp**	**Unlev**	**Lev**	**Other**
Euro-Dollar	0.146	**0.000**	**0.003**	**0.000**	**0.004**	**0.001**	**0.002**	**0.003**
Euro-Yen	0.129	**0.000**	**0.000**	0.828	0.095	**0.000**	**0.008**	**0.003**
Euro-Pound	0.108	**0.005**	**0.001**	0.250	**0.040**	0.285	**0.000**	0.977
Pound-Yen	0.121	**0.010**	**0.000**	0.446	0.452	**0.006**	**0.007**	**0.003**
Pound-Dollar	0.163	**0.000**	**0.000**	0.285	**0.002**	**0.000**	**0.002**	**0.000**
Dollar-Yen	0.169	**0.000**	**0.011**	0.140	**0.040**	**0.003**	**0.007**	**0.003**
Weekly								
Euro-Dollar	0.309	**0.009**	**0.000**	**0.010**	**0.013**	**0.017**	0.071	**0.015**
Euro-Yen	0.328	**0.018**	**0.000**	0.177	**0.019**	**0.000**	**0.000**	**0.000**
Euro-Pound	0.280	0.071	**0.000**	0.057	**0.002**	0.590	**0.000**	0.454
Pound-Yen	0.355	**0.000**	**0.000**	**0.019**	**0.008**	**0.000**	**0.003**	**0.000**
Pound-Dollar	0.335	**0.000**	**0.000**	0.061	**0.004**	**0.000**	**0.000**	**0.015**
Dollar-Yen	0.412	**0.000**	**0.000**	**0.003**	**0.003**	**0.000**	0.452	**0.000**

Table 8

This table presents the probability of informed trading $PIN = \alpha\mu / (\alpha\mu + 2\varepsilon)$ estimated by maximum likelihood for each customer category and for each bilateral currency pair. NC denotes that the maximisation routine failed to converge.

Exchange Rate	Counterparty	Estimated PIN
Euro-Dollar	Corp	0.211
	Unlev	0.365
	Lev	0.316
	Other	NC
Euro-Yen	Corp	0.329
	Unlev	0.479
	Lev	0.353
	Other	0.188
Euro-Pound	Corp	0.092
	Unlev	0.366
	Lev	0.404
	Other	0.142
Pound-Yen	Corp	0.334
	Unlev	0.476
	Lev	0.628
	Other	NC
Pound-Dollar	Corp	0.140
	Unlev	0.291
	Lev	0.433
	Other	0.157
Dollar-Yen	Corp	0.239
	Unlev	0.362
	Lev	0.376
	Other	0.178

Figure 1

This figure presents a stylised passage of order flow through the interbank network. The widths of the arrows denote the size of the order, and the directions of the arrows denote the direction of order flow.

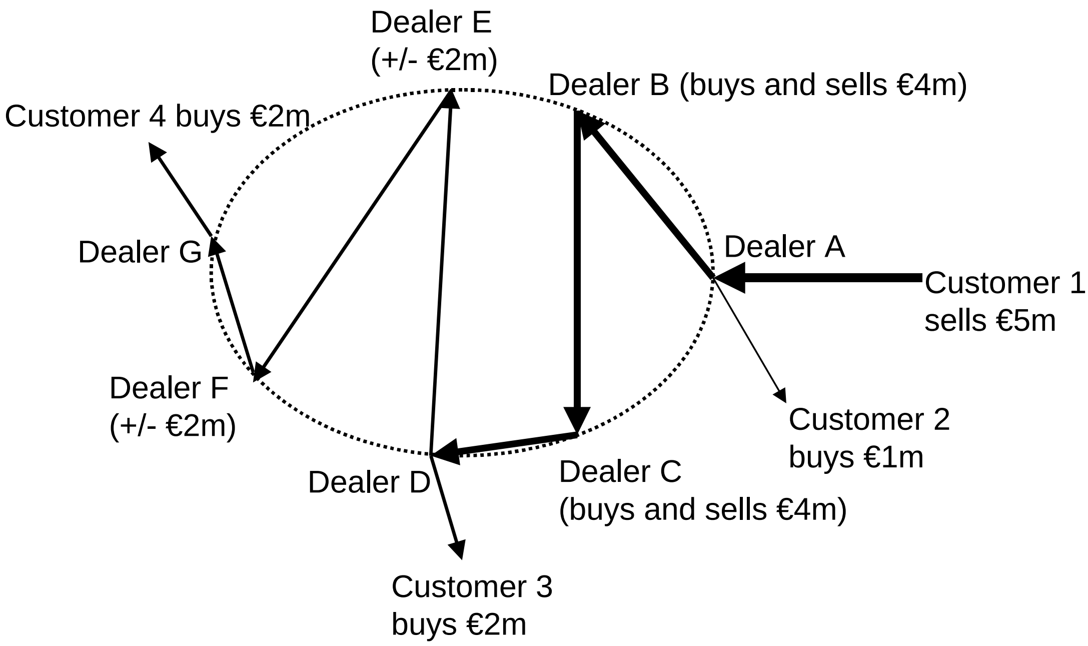

Customer 1 sells €5m to Dealer A for US dollars
Customer 2 buys €1m from Dealer A for US dollars
Dealer A sells €4m to Dealer B for US dollars
Dealer B sells €4m to Dealer C for US dollars
Dealer C sells €4m to Dealer D for US dollars
Customer 3 buys €2m from Dealer D for US dollars
Dealer D sells €2m to Dealer E for US dollars
Dealer E sells €2m to Dealer F for US dollars
Dealer F sells €2m to Dealer G for US dollars
Customer 4 buys €2m from Dealer G for US dollars.

Total customer volume: + €5m + €1m + €2m + €2m = €10m

Total customer order flow: - €5m + €1m + €2m + €2m = €0m

Total inter-dealer volume: + €4m + €4m + €4m + €2m + €2m + €2m = €18m

Total inter-dealer order flow: - €4m - €4m - €4m - €2m - €2m - €2m = - €18m

Dealer A's customer volume: + €5m + €1m = €6m

Dealer A's customer order flow: - €5m + €1m = - €4m