AUSTRALIAN
SCHOOL OF BUSINESS™
THE UNIVERSITY OF NEW SOUTH WALES

Article

Australian Journal of Management
2014, Vol. 39(2) 247–264

DOI: 10.1177/0312896213478435
aum.sagepub.com
SAGE

Measuring the information content of customer foreign exchange orders

Sviatoslav Rosov
School of Slavonic and East European Studies, University College London, UK

F. Douglas Foster
UTS Business School, University of Technology Sydney, Australia

Abstract
This paper investigates whether customer order flow conveys information about future foreign exchange (FX) prices. We use a unique data set from a leading Australian commercial bank that records every FX trade made by the bank in the spot Australian dollar/US dollar market between 2005 and 2010. We find little evidence in support of a cointegrating relation or a statistically significant correlation between customer order flow and FX returns. However, consistent with the liquidity provision role of non-financial customers in Evans and Lyons ((2002) Order flow and exchange rate dynamics. *Journal of Political Economy* 110: 170–180), we find a statistically significant negative correlation between order flow from the diversified economic sector and FX returns. A dynamic analysis suggests that order flow has little or no price impact on FX returns. These results suggest that the non-financial customer order flow of a commercial bank does not carry information about FX prices.

JEL Classification: **G23**

Keywords
Customer orders, foreign exchange, information flows

1. Introduction

Researchers in foreign exchange (FX) market microstructure have, for a number of years, taken as given the idea that information about future exchange rates is conveyed to the market through customer order flow. Customer order flow is the trading activity of a bank's retail customers, as distinct from their corporate counterparties in the interbank market. This mechanism implies that retail FX transactions are somehow correlated with latent macroeconomic variables that determine

Corresponding author:
Sviatoslav Rosov, School of Slavonic and East European Studies, University College London, Gower Street, London WC1E 6BT, UK.
Email: s.rosov@ucl.ac.uk

Final transcript accepted 22 January 2013 by Tom Smith (AE Finance).

exchange rates. While these variables are unobserved, customer order flow is readily measured and in this way the bank's dealers observe noisy signals about future exchange rate movements. This information is subsequently incorporated into the price when the dealers trade in the interbank market.

There have been a number of papers examining various aspects of price formation in the interbank markets. However, few if any papers comprehensively address the question of whether customer order flow is informative about currency prices. This is mainly due to the difficulty in obtaining trading data from commercial banks. Using data from a leading Australian commercial bank, this paper is focused on answering one question: is customer order flow informative about future FX prices?

Much analysis in the literature begins with the Evans and Lyons (2002) trading model, which prescribes a causal relation between customer order flow and the exchange rate. Their three-period trading game sees customers initiate currency demand and supply in the first round of trading before dealers manage the resulting inventory imbalances among themselves in the second round. Any remaining imbalances are passed back to the customers in the third round of trading with the help of compensating premiums. The testable implication is a cointegrating relation between cumulative order flow and the exchange rate in levels.

We use the Gereben et al. (2006) interpretation of the model where first- and third-round customers are different types of customers. First-round or push customers are financial customers that potentially trade on information. Third-round or pull customers are liquidity providers (typically diversified customers) drawn into the market by the premium offered by the dealer.

Our results do not provide evidence in support of a cointegrating relation between cumulative daily aggregate order flow and the exchange rate. Nor is there a cointegrating vector between individual economic sector cumulative flows (i.e. the diversified sector) and the exchange rate. These conclusions remain when a weekly or monthly frequency is used or when sub-periods are analysed.

The lack of a cointegrating vector motivates a static analysis of the relation between exchange rate returns and net order flow. This gives qualified support to the Evans and Lyons (2002) model. Firstly, there is no relation between aggregate order flow and the exchange rate, consistent with the predictions of the model. Further, we observe a significantly negative relation between third-round order flow made by pull customers (diversified sector) and the exchange rate, also consistent with the model. However, there is limited evidence in support of a significant, positive relation between first-round order flow made by push customers (financial sector) and the exchange rate. A sub-period analysis does not change these conclusions, but does suggest that trading during the 2008 financial crisis sub-period is driven by factors exogenous to the three- period model (or any standard microstructure mechanisms).

Next, we fit a Vector Autoregression (VAR) model to find a dynamic relation between customer order flow and FX returns. The VAR equations can be interpreted as the result of a structural trading model, as in Hasbrouck (1991a). In this case the impulse response of the VAR system gives the researcher the long-run price impact of the order flow. The variance decomposition is interpreted as the proportion of total return information that is conveyed through order flow (see Hasbrouck, 1991b). The VAR results show no systematic statistical link between customer order flow and FX returns at any frequency. The variance decompositions suggest that customer order flow has a statistically significant double-digit information share at a monthly frequency (but not at daily and weekly frequencies).

The contribution of this paper is in the unique customer order flow data set that has a level of detail and length of sample period that is, as far as we know, unmatched in the existing literature. Using a variety of techniques does not identify any significant relation between the order flow of the bank's retail customers and the Australian dollar (AUD)/US dollar (USD) exchange rate. The data may not be a representative sample of market activity due to the unidirectional trades of customers and a relative lack of large institutional customers.

The remainder of this paper is structured as follows. Section 2 presents an introduction to FX microstructure. Next, Section 3 examines the data set. The results in Section 4 begin with a synopsis of the Evans and Lyons (2002) model in Section 4.1, whose testable implications are investigated in Section 4.2. A static regression analysis is conducted in Section 4.3 and is followed by a dynamic VAR analysis in Section 4.4. Section 5 concludes.

2. Background

2.1. Institutional

The global FX market is the largest market in the world. Combined, spot and forward trading amounts to US$1.4 trillion per day (BIS, 2007). Including swaps, this figure rises to US$3.2 trillion. The FX market is largely unregulated with a decentralized market structure where participants are physically separated from one another and transactions take place over communications networks.

There are two tiers in the FX market. The first tier sees trading and currency demand from customers. The trading in this market is almost completely opaque, with customer trades unknown outside individual banks; aggregate figures for customer trading volume are published only triennially by the Bank for International Settlements (e.g. BIS, 2007). There is commonly recognized heterogeneity within customer trading. Financial customers account for about 40% of trading volume (BIS, 2007) and include Central Banks, banks, mutual funds, leveraged investors, etc. These customers usually demand FX for speculative or investing purposes – they are interested in currency as a store of value (Osler, 2008). Corporate customers trade currency as part of their day-to-day business operations (e.g. exporting minerals). These customers are interested in currency as a medium of exchange and do not tend to divert resources away from their core business to speculate on price movements (Osler, 2008). Corporate trading accounts for about 20% of customer trading volume.

The second tier of the market is interdealer trading. This accounts for about 38% of trading (BIS, 2007). The interdealer market is semi-opaque in the sense that it is possible to observe quotes from other dealers and infer activity, but full information about the trades is not available. The interdealer market is virtually unrecognizable from its form in the 1990s and earlier. Direct interdealer trading between banks is almost non-existent (Osler, 2008), with trades happening almost exclusively through electronic communication networks (ECNs) such as the Reuters Dealing systems or Electronic Broking Services (EBS).

It is important to note that this separation of a customer market and an interdealer market does not exist in the equity microstructure world, yet it is vital in FX where most of the trading volume occurs in the interdealer market. Despite this, all currency demand is, ultimately, driven by customers. Retail and corporate customers looking to trade currency with the bank will be dealt with by the sales team. The sales team has no access to the interbank market and serves simply as a conduit between customers and dealers – both passively receiving orders as well as actively soliciting trades from its address book of customers. The dealers are the agents who must manage a bank's position in a particular currency.

For the majority of trades, the sales team will quote and consummate a trade with a customer without the dealer's knowledge. Since most trades are with established customers of the bank and for relatively small amounts, this system prevents dealers from being overwhelmed with routine trades. Further, established customers can trade currency at any time, up to an arranged credit limit, through an online portal that most banks operate. Many customer trades therefore bypass even the sales team.

Once a trade is consummated by the sales desk, the exposure is electronically routed to the dealer managing the bank's position in that currency. Most trades thus appear in the dealer's inventory as exogenous shocks in the tradition of the inventory model literature. For these sorts of trades, the dealer neither knows nor cares about the identity of the counterparty.

The rather mechanical process described above breaks down for very large trades, which are dealt with in a far more cooperative manner. There are two considerations: firstly, the bank wants to do business with customers that make large trades; and, secondly, it is no longer acceptable for the dealer to be surprised with a large inventory shock. It is likely that the customer is shopping around other banks so a race is on to quote a favourable price at which the customer will trade. In addition, because a large trade would have a large impact on a dealer's inventory, the relevant dealer is advised of any pending large trades. This allows the dealer time to build up an inventory buffer and/or scan the market in order to fill the order at the best price.

Customer trading causes the bank to experience inventory shocks. The responsibility of managing these shocks is given to dealers. Dealers are players that are motivated by profit, with large proportions of their remuneration coming from profit-related bonus payments. In pursuing profit, they are constrained by position and loss limits in order to manage risk. In managing inventory, dealers are commonly assumed to end each trading day flat (i.e. with no inventory position) (Osler, 2008). It should be noted that with 24-hour trading, the definition of a trading day is uncertain. Discussions with practitioners suggest that dealers do not have to end each day exactly flat, since inventory imbalances can be passed onto the dealers in the next time zone. Nevertheless, there is a concerted effort to keep these imbalances to a minimum, particularly on Fridays in preparation for the weekend (Bjønnes and Rime, 2005). Finally, dealers manage inventory very intensively with the half-life of inventory positions ranging between 5 and 30 minutes (Bjønnes and Rime, 2005).

Inventory is managed through interdealer trading (Osler, 2008) in preference to price shading (i.e. waiting for and/or encouraging reversing customer trades). The extremely liquid interbank market allows dealers to off-load inventory shocks quickly, cheaply and easily, rendering the practice of price shading redundant. In this market, price shading would provide a clear and ultimately costly signal of an inventory imbalance.

While inventory management is the most important task of dealers, they also speculate and seek out arbitrage opportunities. Cheung and Chinn (2001) find that dealers make a large part of their profit by speculating on rate movement, not spread. This is opposed to the bank as a whole, which makes most money from charging spreads on customer trades.

2.2. FX microstructure models

Relative to most microstructure models, the first unique characteristic of FX trading is that the dealer may know the identity of the customer making a trade. Traditionally, the market makers have not been able to tell informed from uninformed; in fact, inferring who is informed has been the key objective of the market maker. However, over time FX market makers are likely to have a noisy signal of who is informed – they might pay more attention to an order from a large corporate entity or a large bank than from a small retailer. A second challenge is that few of the traditional information models (as distinct from inventory models) make any mention of the market maker's inventory position. However, there is a lot of evidence that FX traders manage their inventory intensively so it is not realistic to assume this problem away. A desirable FX model should combine the information and inventory effects.

Evans and Lyons (2002) were the first to present an empirical model of FX rates with significant explanatory power. Their microstructure approach used interdealer order flow to explain FX rate movements and provided a successful adaptation of the equity microstructure models to the problem of FX markets. Their central empirical finding was that their model of daily Deutsche Mark (DEM)/USD log returns achieved a high R^2 statistic and implied that US$1 billion of positive order flow increases the DEM price of a US dollar by 0.5 percent.

Bjønnes et al. (2005) provide an excellent overview of the model, which is reproduced in short form here. At the start of the first period, dealers receive public macro information and quote a common price (since all dealers have the same information). Dealers then receive private customer flow that is not observed by the market. Customer order flow creates inventory imbalances for the dealers, so in the second round they trade among themselves to share inventory risk and speculate on their private information. They again quote identical prices, since they do not want to reveal their private information. At the end of round 2, total net interdealer order flow is displayed to all participants. The net interdealer trades initiated by a dealer in round 2 are proportional to the level of her customer order flow in round 1. However, this interdealer trading does not necessarily result in the desired inventory position, since other dealers are also off-loading their positions. Dealers induce retail customers to absorb excess inventory in a third round of trading. These customers need to be compensated to hold a non-optimal level of currency.

Based on this model, Bjønnes et al. (2005) argue that you would expect to see a positive correlation between round 1 excess demand for a currency and the value of the currency. Since excess positive demand cannot be entirely cleared in the interbank trading round, a positive premium must therefore be paid (i.e. currency appreciation) to induce customers to re-absorb some of this demand in period 3. Dealers use information on net interdealer order flow in round 2 to set this price. From the total net interdealer order flow the dealers deduce the total amount that the public needs to absorb. The dealers also know the public's risk-bearing capacity so the common round 3 price is at a premium relative to the round 2 price. For example, if customers buy AUD in round 1 then the aggregate interdealer order flow at the end of round 2 will be positive as dealers buy back AUD. Since dealers lay off all their inventories during round 3, this means that the aggregate customer orders in round 3 should be of a similar size but opposite sign to those in the first round.

2.2.1. Customer heterogeneity. The literature on FX microstructure can be put into several distinct categories. Firstly, there are the papers that replicate the results of Evans and Lyons (2002) while contributing new econometric techniques, unique data sets or introducing new issues, such as feedback trading. For example, Berger et al. (2005), having the advantage of a much longer five-year data set from EBS, find a strong positive association between order flow and returns at frequencies between one minute and one day. They estimate a similar impact of an extra US$1 billion in order flow to Evans and Lyons (2002). Specifically, it should raise the Euro (EUR) price of USDs by around 0.4 percent.

Another set of papers examines customer trading data and investigates the importance of customer heterogeneity. Bjønnes et al. (2005) argue that since market-making banks try to end each day flat, it is not plausible that they are the main overnight liquidity providers in the FX market. Bjønnes et al. (2005) suggest candidates for liquidity-demanding and liquidity-supplying customers. They put forward the liquidity hypothesis to explain the correlation between order flow and exchange rates. Specifically, a typical liquidity-demanding customer (i.e. a round 1 customer in the Evans and Lyons (2002) model) is an aggressive financial institution, the so-called push customer.

This customer is active in the sense of approaching the market maker for exogenous (to the model) reasons – they are a liquidity demander. A typical liquidity-supplying customer (round 3 customer), the so-called pull customer, is more likely to be a non-financial customer that is attracted into the market by the risk premium offered by the market maker to absorb the excess inventory at the end of the trading period.

Therefore, order flow from a push trader should be positively correlated with contemporaneous price changes. However, flows of pull customers should be negatively correlated with contemporaneous price changes. Ignoring this distinction has the potential to wash out any significant effects that would otherwise be observed across customers. Bjønnes et al. (2005) outline testable empirical implications of the Evans and Lyons (2002) model. These are that there should be a cointegrating relation between the level of the exchange rate and the first-period order flow (similar cointegrating relations should exist with the second or interdealer and third-period order flow); and third-round liquidity supply should net out first-round liquidity demand.

Consistent with their hypothesis, Bjønnes et al. (2005) present two pieces of evidence suggesting that non-financial customers are the main liquidity providers in the overnight FX market. Firstly, the net position of non-financial customers is negatively correlated with the FX rate (financial customers are positively correlated). Secondly, changes in the net position of non-financial customers are forecasted (Granger-caused) by changes in the net position of financial customers (suggesting that non-financials are passive players, consistent with liquidity provision).

This hypothesis is consistent with the results of Carpenter and Wang (2003), who find that the central bank has the greatest price impact, followed by non-bank financial firms (e.g. hedge funds or mutual funds), with non-financial corporations having the least price impact. Marsh and O'Rourke (2005) similarly find that order flow from non-financial corporations is negatively correlated with FX returns, while order flow from financial companies is positively correlated with FX returns.

While the relation between order flow and returns for different customers is typically examined using cointegration (Osler, 2008), Gereben et al. (2006) provide a novel and complimentary approach that they argue is consistent with Kyle (1985), where pull customers and dealers jointly act like the market maker. The dealers set prices and the pull customers are the ones who are willing to take positions and provide liquidity. The authors find that foreign order flow (comprised of large financial firms) helps explain exchange rate fluctuations. However, domestic bank and non-bank order flow (pull customers) are insignificant. Here the authors likely have a geographical effect that is a proxy for the financial/non-financial dichotomy found in other papers.

2.2.2. Information hypothesis. A relatively recent branch of the literature has begun to address the question of whether the relation between order flow and FX returns is due to an information effect. Osler (2008) argues that the majority of evidence supporting the information hypothesis is based on the permanence of the exchange rate response to order flow. Inventory effects imply transient price impacts (to solicit unwinding trades) and are not consistent with the long-run impacts observed in the literature. A permanent impact is consistent with the liquidity hypothesis if shifts in liquidity demand or supply are permanent (Osler, 2008). However, a permanent relation is inevitable if order flow carries private information.

Payne (2003) is the first paper in FX microstructure to take advantage of the Hasbrouck (1991a) VAR approach. The VAR model allows a dynamic analysis of the relation between order flow and returns, which in turn enables researchers to examine the persistence of the price impact of order flow using impulse response functions (IRFs). In addition, it is possible to decompose the variance of returns into components caused by returns and order flow. The latter is argued by Hasbrouck (1991a) to be an indicator of the information content of order flow.

Using Reuters data on the USD/DEM exchange rate, Payne (2003) finds that an unexpected market buy (only the sign of the trade is used as an explanatory variable, not the size), leads to a permanent quote revision of 1 pip (i.e. DEM 0.0001). Further, he estimates that 40% of all permanent price variation is due to order flow. Similar results suggesting the permanence of the price impact of order flow are found in Killeen et al. (2006) and Berger et al. (2005).

The information hypothesis posits that order flow affects exchange rates because it carries private information. This private information may be in the form of unobservable fundamental determinants of the exchange rate, which are compounded in each dealer's customer order flow (as modelled by Evans and Lyons, 2005). Customer order flow therefore gives the dealer a signal about the future values of these fundamental determinants. This information is incorporated slowly into the exchange rate through interdealer trading. Reitz et al. (2007) argue that evidence of such a long-run equilibrium relation between customer order flow and exchange rates can be tested by searching for cointegration.

3. Data

The data we use in this study covers the period from February 2005 to April 2010, comprises 6,478,758 line items and contains information on anonymous ID code, industry description, type of trade (internal or external, spot/forward/swap) and trade details (currency pair, amount in both currencies, date and time of trade, broker, dealer information, trade identification stamp). Compared to prior work, this data is unique because we are able to track individual customers through the sample period with the aid of an anonymous ID number and assign each customer to an economic sector and a specific industry within that sector.

Trades within the bank are classified as either spot, outright forwards or swap trades. A spot trade has two days to maturity (this is how long it takes to clear a spot trade). Trades maturing the next day are called overnight trades, trades with one day to maturity are called 'tom-next' and trades with three days to maturity are called 'spot-next'. Thereafter, trades are known as outright forwards. We define a spot trade to be any overnight, tom-next or spot trade. The breakdown of trades with customers is as follows: spot trades (44% of all trades), swap trades (43%) and outright forwards (15%).[1]

The most heavily traded currency pair in the customer data is AUD/USD (58% of spot trades), followed by AUD/EUR (14%) and AUD/Great British Pound (GBP; 7%). Although there are hundreds of distinct currency pairs, the top five command 85% of the total number of customer trades. Anecdotally, the vast majority of the bank's customers are considered to be small to medium sized importers that exchange AUD for USD (and to a lesser extent the EUR) in order to pay for imports into Australia. Therefore, our analysis will focus on customer spot AUD/USD trading.

3.1. Customers

Each customer is assigned to one of 11 economic sectors. Table 1 shows that the bank's retail customers are mostly diversified firms. There are 14,030 customers in total and during the five years of our sample they trade an average of 116 times. Approximately 40% of retail customers come from the diversified industry, with financial, agricultural and transport firms also being quite common. The distribution of trades among these customers is approximately uniform in the sense that no single customer accounts for more than 1% of all trades. However, the median number of trades for diversified and financial firms is far smaller than the average, suggesting that only a subset of these customers accounts for the majority of trades by that industry group. The opposite is true for

Table 1. Customers.
This table shows (by economic sector) the number of distinct retail customers. Also shown are the average and median number of trades per customer in each sector and also the total number of trades made by all customers in that sector.

Sector	No. of customers	Avg. no. trades/ customer	Median no. trades/customer	Total sector trades
Diversified	5802	154	11	895,927
Financial	2705	118	1	319,939
Agriculture	603	237	1746	142,953
Transport	543	171	1	92,770
Missing	3013	16	1	47,356
Resources	245	157	5	38,453
Media	141	174	4	24,589
Energy	175	119	1093	20,822
Government	169	112	2	18,987
Property	372	48	23	18,005
Unallocated	262	27	1	7198

agriculture and energy firms, with customers either trading very frequently or not at all. This causes the median number of trades to be higher than the average.

The diversified and financial economic sectors dominate trading with the bank. They are first and second most active, respectively, in spot trades, together accounting for 76% of spot AUD/USD trades. The other economic sectors have single-digit shares of bank trading activity.

3.2. Trading activity

Daily trading activity follows an inverted 'W' pattern where the trough represents the lunch break. We define the trading day as 7:30 a.m.–6:30 p.m. local Sydney time, Monday–Friday. Consistent with Treepongkaruna et al. (2012) (who build on Treepongkaruna and Gray, 2009), we exclude trading that occurs over the weekend. Trading falls overwhelmingly inside this working day for the spot, swap and forward markets. The inter-trade time statistics (not shown) show that the median inter-trade time is between 1 and 3 minutes for customer trades (depending on the trade type).

Discussions with traders suggest that a daily frequency is approximately the hurdle at which they would expect the signal-to-noise ratio to become acceptable for customer data. The average time between trades for the customers of the Bank is almost 83 days across 7178 customers. The median is 34 days with the 5th and 95th percentiles being 4 and 319 days, respectively. Hence, it seems that the minimum frequency we should use for the customer VAR analysis is weekly or even monthly.

3.3. FX rate movements

Figure 1 shows the AUD/USD exchange rate over our sample period from February 2005 to December 2009. There was a large movement of the AUD/USD exchange rate during the global financial crisis in 2008. Over the entire sample period the AUD/USD exchange rate fell from near parity to 60 US cents per dollar and returned to near parity. This suggests a sub-period analysis, where the sub-periods are delineated by vertical lines in Figure 1 and correspond to the following: UP1 (24 February 2005–17 June 2008), DOWN (18 June 2008–13 April 2009) and UP2 (14 April 2009–9 December 2009).

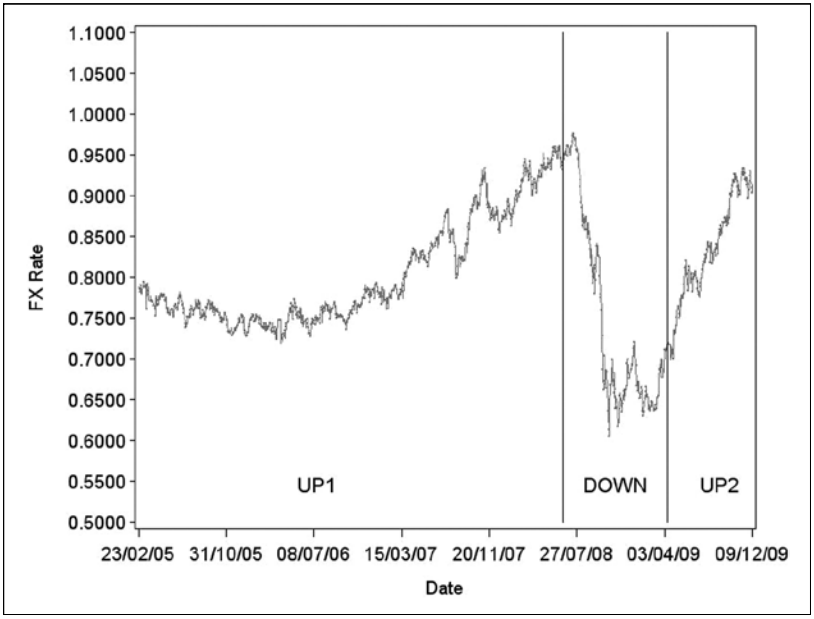

Figure 1. Australian dollar (AUD)/US (USD) dollar foreign exchange (FX) rate over time.
This figure shows a time series plot of the AUD/USD exchange rate. The plot is divided into three sub-periods – UP1 (24 February 2005–17 June 2008), DOWN (18 June 2008–13 April 2009) and UP2 (14 April 2009–9 December 2009).

3.4. Order flow

Order flow in our data is defined so that a positive number is a buy order initiated by the customer. Due to the possibility of varying spreads, we use the spread mid-quote as the FX rate. The order flow variable suffers from a scaling problem, with orders varying greatly in size (descriptive statistics not shown), so there will be influential data points. One approach is to take the logarithm of order flow, but this is not possible since order flow takes negative values. A common approach in microstructure is to simply use a trade indicator variable. We have two trade indicator variables. The first, signflow, takes the value +1 for a buyer-initiated purchase and –1 for a seller-initiated sale. It can also take a value of 0 if during that clock-time period there is no trade. Since there are many trades in each clock-time period, signflow measures whether there was net buying or net selling pressure during that period.

The second variable, tradeflow, is similar to signflow in that it measures net buying or net selling pressure during a particular period; however, rather than taking the value +1 or –1 it counts the net number of purchases or sales during that period. For example, if during some period there were 15 purchases and 7 sales, the signflow variable would take the value +1 while the tradeflow variable would take the value +8. Finally, like signflow, tradeflow (and order flow) takes the value 0 if there are no trades during a particular clock-time period.

4. Results

In this section we present the results from three approaches to searching for information content in spot customer trades. Section 4.1 tests for cointegration between customer order flow and the exchange rate. This is followed in Section 4.2 with a static regression analysis and in Section 4.3 with a dynamic VAR analysis.

4.1. Cointegration

There are a handful of papers that use cointegration procedures in the microstructure literature. Killeen et al. (2006) and Bjønnes et al. (2005) both use Johansen cointegration procedures to test the implications of the Evans and Lyons (2002) model.

Bjønnes et al. (2005) test for the first- and third-round customer order flow cointegrating relation and find evidence supporting the Evans and Lyons (2002) model. Bjønnes et al. (2005) show that non-financial customer orders are negatively cointegrated with the exchange rate, whereas financial customer orders are positively cointegrated. They also find that financial order flow Granger-causes non-financial order flow (but not vice versa), which is consistent with non-financial customers being passive liquidity providers. Killeen et al. (2006) test for cointegrating relations with inter-bank data and find a long-run relation between cumulative order flow and the exchange rate, suggesting at least a portion of the price impact of order flow is permanent. This is consistent with order flow conveying information, since information effects should be permanent (Hasbrouck, 1991a).

We test for cointegration by first determining whether a cointegrating vector exists between the non-stationary price and cumulative order flow time series. We use the Engle–Granger test since we have only two variables so, therefore, there is at most one cointegrating vector.

The results suggest that the residuals from a linear regression of the exchange rate on cumulative order flow are non-stationary, which implies that the exchange rate and cumulative order flow are not cointegrated. This result is the same whether the Augmented Dickey-Fuller (ADF), Phillips-Perron (PP) and Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test is performed. Using the Johansen procedure to test for cointegration does not change the findings. We have also searched for cointegration at the lower weekly and monthly frequencies. In both cases the results are qualitatively the same as for the daily frequency, with no evidence of cointegration by either the Engle–Granger or Johansen procedures.

Because cointegration is evidence of a long-run relation between two non-stationary variables, it is possible that the financial crisis of 2008 is clouding the cointegration tests. The Johansen procedure overwhelmingly rejects the existence of any cointegrating relation in any sub-period. We also use the Johansen procedure to test for evidence of cointegration using the trade indicator variables signflow and tradeflow (this is also done by Reitz et al., 2007). The results for both signflow and tradeflow are qualitatively identical to those for order flow (including for sub-periods) and are not reported.

The weight of evidence suggests there is no cointegrating relation between the exchange rate and cumulative order flow, which is contrary to the findings of Bjønnes et al. (2005). In turn, this means there is no error correction representation between order flow and the exchange rate and, therefore, we estimate a relation in first differences.

4.2. Regression

The Evans and Lyons (2002) model does not have explicit testable implications on the relation between exchange rate returns and order flow. Nevertheless, a simple linear regression between FX returns and order flow is commonly employed (see, for example: Gereben et al., 2006). Such a regression is used by Bjønnes et al. (2005) as additional analysis beyond their cointegration tests.

They argue that while the cointegration analysis establishes evidence of a long-run relation between order flow and the exchange rate, ordinary least squares (OLS) regression on first differences sheds light on short-run dynamics.

Gereben et al. (2006) find a significant relation between order flow and FX returns at a daily frequency. This relation is positive for foreign customer order flow and negative for domestic customer order flow. The authors argue that this result is consistent with the push (financial)/pull (non-financial) dichotomy proposed by Bjønnes et al. (2005) because, in the Hungarian context, financial firms tend to be foreign customers while domestic Hungarian customers tend to be non-financial firms.

Bjønnes et al. (2005) also estimate a regression between pull customer order flow and FX returns. They argue that because the original cointegration relation is equivalent for first- and third-round order flow (only the sign changes to indicate whether liquidity is being demanded or supplied), the inference is invariant to the type of order flow used – first-round push order flow or third-round pull order flow. The authors include only one order flow variable (either financial or non-financial), because including both would give rise to multicollinearity. They find that the coefficient on non-financial customer order flow is negative and almost exactly the opposite of the positive coefficient on financial customers. This result is invariant to the frequency of data used, ranging from 5 to 90 days. Bjønnes et al. (2005) argue that this shows non-financial customers provide liquidity at shorter horizons.

An interesting result in Bjønnes et al. (2005) is that at a daily frequency, non-financial customers are also push customers, having a positive relation to FX returns. This switches to a negative relation at a five-day frequency. The coefficient on market-making bank order flow is negative at all frequencies but decreases in importance beyond a five-day frequency. The authors conclude that this evidence suggests that market-making banks act as liquidity providers at the daily horizon, with non-financial customers becoming more important as the time horizon increases.

Table 2 shows results of regressions of returns on total or aggregate order flow (odd columns) and economic sector order flow (even columns) for each of the three trade variable definitions. Total customer order flow does not appear to be related to FX returns, which is consistent with the Evans and Lyons (2002) three-period model. This model predicts that net bank-wide aggregate customer order flow should only be randomly different from zero and should not be correlated with the exchange rate (Gereben et al., 2006).

There is, however, a relation when the signflow and tradeflow variables are used, with the coefficient on the aggregate flow variable being significantly negative in these two tests. This is likely because ignoring the size of the trade from the flow variable increases the importance of the diversified sector trades (which are numerous but small in dollar terms) relative to financial sector trades (which are less numerous but far larger in dollar terms).

The results for economic-sector-based regressions are also consistent with push–pull customer segregation. We see that for all trade variables, the coefficient on the push customer (financials) is positive, albeit insignificant. The coefficient on the diversified firms (pull customers) is negative and significant for all trade variables. This supports the findings of Gereben et al. (2006), who show that the order flow of non-informed liquidity providers (i.e. third-round pull customers) correlates negatively with the exchange rate.

The miscellaneous group also has a significant negative coefficient. It is likely that for the purposes of our analysis, the diversified and miscellaneous economic sectors could be bundled into a broad non-financial sector, as seen in other research. The results for the government order flow are less clear, with insignificantly negative coefficients using order flow and signflow, but a significantly positive coefficient using tradeflow. It should be noted that government customers do not include central banks, which are included in financials.[2]

Table 2. Regression analysis – daily.
The table reports the regression results using mid-quote daily returns as the dependent variable. The independent variables are comprised of various definitions and sub-samples of flow data. Three trade variables are used to define the flow data: order flow, signflow and tradeflow. Signflow takes the value +1 for a buyer-initiated purchase and –1 for a seller-initiated sale (0 if during that clock-time period there is no trade). Tradeflow counts the net number of purchases or sales during that period (0 if no trades during that clock-time period). For example, if during some period there were 15 purchases and 7 sales, the signflow variable would take the value +1, while the tradeflow variable would take the value +8. Two specifications of regressions are estimated. Columns 1, 3 and 5 have an aggregate flow variable as the independent variable. Columns 2, 4 and 6 have flow variables split by economic sector. *t*-statistics are shown below each estimate in parentheses. *, ** and *** show significance at the 10%, 5% and 1% levels, respectively.

	Order flow		Signflow		Tradeflow	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.0001	−0.0002	−0.0002	−0.0006*	−0.0040***	−0.0040***
	(0.4300)	(–0.5700)	(–0.5800)	(–1.7300)	(–6.5300)	(–6.3400)
Total	0.0000		−0.0009***		−0.0000***	
	(–0.3400)		(–3.1600)		(–7.5000)	
Financial		0.0000*		0.0004		0.0000
		(1.6800)		(1.4900)		(0.8800)
Diversified		−0.0001***		−0.0018***		−0.0000***
		(–4.9900)		(–5.6000)		(–4.1200)
Government		−0.0001		0.0000		0.0002***
		(–1.3600)		(–0.1600)		(2.7000)
Miscellaneous		−0.0000***		−0.0009***		−0.0000***
		(–3.8600)		(–3.3300)		(–3.0000)
N	1510	1510	1510	1510	1510	1510
Adj. R^2	−0.0006	0.0251	0.0059	0.0274	0.0353	0.0414

Table 3 lists results for our regression analysis at a monthly frequency. While the inference for total aggregate order flow remains the same, results using economic sector order flows are different from those at a daily or weekly frequency. There does not appear to be a meaningful relation between sector flows and the exchange rate at a monthly frequency. Table 3 shows no significant relation between diversified flows and the exchange rate using any trade variable, despite there still being a significant coefficient on total flows for the signflow and tradeflow specifications. Hence, monthly frequencies appear to be too coarse to capture the liquidity provision outlined by Evans and Lyons (2002).

4.3. Sub-period analysis

Gereben et al. (2006) find that central bank interventions are highly informative about the exchange rate during times of market turmoil, but had no measurable impact during calmer trading periods. Our sample period has a period of very large market turmoil where the AUD fell from near parity to around 0.65 AUD/USD. However, customer order flow (in this case largely non-central bank orders) may not be informative, since trading might be driven by uncertainty and the need to unwind trading positions. In this section we repeat the regression analysis on the three sub-periods defined previously. The regression results by sub-period are shown in Table 4.

The regressions are run only for the order flow trade variable. Given that the UP1 period comprises most of the sample period, it is not surprising to see that the results mirror those of the full

Table 3. Regression analysis – monthly.
The table reports the regression results using mid-quote daily returns as the dependent variable. The independent variables are comprised of various definitions and sub-samples of flow data. Three trade variables are used to define the flow data: order flow, signflow and tradeflow. Signflow takes the value +1 for a buyer-initiated purchase and –1 for a seller-initiated sale (0 if during that clock-time period there is no trade). Tradeflow counts the net number of purchases or sales during that period (0 if no trades during that clock-time period). For example, if during some period there were 15 purchases and 7 sales, the signflow variable would take the value +1, while the tradeflow variable would take the value +8. Two specifications of regressions are estimated. Columns 1, 3 and 5 have an aggregate flow variable as the independent variable. Columns 2, 4 and 6 have flow variables split by economic sector. *t*-statistics are shown below each estimate in parentheses. *, ** and *** show significance at the 10%, 5% and 1% levels, respectively.

	Order flow		Signflow		Tradeflow	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	−0.0029	−0.0063	−0.0257***	−0.0345	−0.0674***	−0.0606***
	(–0.4200)	(–0.5900)	(–2.8800)	(–1.5100)	(–3.1600)	(–2.7400)
Total	0.0000		−0.0335***		−0.0000***	
	(–1.2100)		(–3.7500)		(–3.3800)	
Financial		0.0000		−0.0486***		0.0000
		(–0.3800)		(–4.3000)		(–0.7000)
Diversified		−0.0001		0.0104		0.0000
		(–1.2000)		(0.5400)		(–0.7800)
Government		0.0000		−0.0041		−0.0000
		(–0.0700)		(–0.6400)		(–0.1400)
Miscellaneous		−0.0000*		0.0005		−0.0000**
		(–1.7900)		(0.1000)		(–1.9900)
N	70	70	70	70	70	70
Adj. R^2	0.0066	0.0096	0.1595	0.1929	0.1315	0.0994

analysis. That is, there is no relation between total aggregate order flow and the exchange rate. However, there is a significant negative relation between the (presumed) third-round pull order flow from diversified customers and the exchange rate.

We expect little information content in trades during the DOWN period if trading during this time was driven by sentiment and a global tightening of liquidity. It is therefore surprising to see a significantly negative relation between total order flow and the exchange rate. It is not immediately clear why this should be the case.

The UP2 sample period is, as predicted, similar to the UP1 sub-period, with total order flow being uncorrelated with the exchange rate. However, there is now a significantly positive relation between financial order flow and the exchange rate, lending support to the hypothesis that they are first-round push customers. Consistent with this push–pull breakdown is the usual statistically significant negative coefficient on the diversified order flow. The daily sub-period regression analysis suggests that something unusual (relative to the UP1 and UP2 periods) is happening in the DOWN sub-period. Sub-period differences are not clearer at a weekly frequency and are not shown.[3] A monthly sub-period analysis is not possible, because there are too few observations to be able to make an inference during the DOWN and UP2 sub-periods.[4]

While this section presents an analysis of the static contemporaneous relation between order flow and FX returns, the next section presents a dynamic analysis using VAR techniques.

Table 4. Regression analysis – daily order flow by sub-period.
Shown below is the sub-period regression output using mid-quote daily returns as the dependent variable. The independent variables are various sub-samples of order flow data. The sub-periods are defined as follows: UP1 (24 February 2005–17 June 2008), DOWN (18 June 2008–13 April 2009) and UP2 (14 April 2009–9 December 2009). Two specifications of regressions are estimated. Columns 1, 3 and 5 have aggregate order flow as the independent variable. Columns 2, 4 and 6 have order flow variables split by economic sector. *t*-statistics are shown below each estimate in parentheses. *, ** and *** show significance at the 10%, 5% and 1% levels, respectively.

	UP1		DOWN		UP2	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.0002	0.0000	−0.0017	−0.0009	0.0008	−0.0010*
	(0.7200)	(–0.0600)	(–1.4300)	(–0.7300)	(1.5300)	(–1.7400)
Total	0.0000		−0.0000**		0.0000	
	(0.1600)		(–2.1400)		(0.7600)	
Financial		0.0000		0.0000		0.0000**
		(1.2100)		(–0.4700)		(2.0400)
Diversified		−0.0000***		0.0000		−0.0002***
		(–3.8300)		(–0.7900)		(–5.6500)
Government		0.0000		0.0001		−0.0005***
		(–0.8100)		(0.3500)		(–3.0000)
Miscellaneous		0.0000		−0.0001***		0.0000
		(–0.9900)		(–3.3300)		(–1.3200)
N	863	863	214	214	433	433
Adj. R^2	−0.0011	0.0148	0.0165	0.0386	−0.001	0.0935

4.4. VAR

Hasbrouck (1991a) introduces VAR analysis in microstructure research to understand the information content of order flow. Short-run price impact is interesting because it describes how an order affects price upon its submission. Hasbrouck (1991a) argues that to determine the information content of order flow, the long-run price impact of the order is also important. This is given by the IRF. Further, the Forecast Error Variance Decomposition (FEVD) allows a researcher to decompose the forecast error variance into components that arise from uncertainty in a particular variable. In the literature, this is interpreted as the proportion of price variation that is explained by the information arising from order flow.

There are several papers that perform a VAR analysis on exchange rate data. However, they mostly use interbank data, while we concentrate on customer orders. For example, Payne (2003) finds that asymmetric information accounts for around 60% of average bid–ask spreads in the USD/DEM market. In addition, he finds that 40% of all permanent price variation (the FEVD) is due to order flow. Berger et al. (2005) use five years of brokered EBS data and estimate the long-run response to a one standard deviation shock (19.2 million EUR) is around 1.2 basis points. Their variance decomposition results are similar to Payne (2003), with the share of variance attributable to order flow shocks being around 40%.

We run a two-equation system using FX returns and the signflow and tradeflow variables where order flow logically precedes returns. It is also possible to incorporate a contemporaneous returns term. Danielsson and Love (2004) argue that feedback trading is an inevitable consequence of time aggregation of order flow models. Their results suggest that at the 1 minute frequency, the feedback trading parameter is positive and significant and the IRF following an order flow shock is larger when feedback trading is incorporated.

However, the incorporation of feedback trading makes the VAR system non-identified and requires the use of (likely) error-inducing instrumental variables to identify structural parameters. The addition of additional equations, such as inventory or trader types, poses difficulties both in terms of data availability, as well as the logical ordering of the Cholesky decomposition. Having only two equations alleviates this problem, since one must only decide whether the flow variable precedes price revisions or vice versa. We follow the practice of other researchers in having the flow variable logically precede price revisions.

Table 5 reports VAR estimation results for customer order flow data. The returns are, as before, calculated using the approximated mid-quote price series. The trade indicator variables signflow and tradeflow are used to represent orders (to avoid the skewness of order flow). The first statistic presented is the long-run price impact for a model with five lags of each variable at a daily frequency, suggesting that past prices and orders are relevant for one business week. This is consistent with dealers managing positions over the weekend. For a weekly analysis, four lags (one month) are used, while the monthly analysis is limited by data availability and thus we use only two lags.[5]

The second statistic presented in Table 5 is the FEVD of returns. This is the proportion of total FEVD associated with the order flow innovation. The intuition here is that private information enters prices through unexpected order flow. Therefore, the FEVD associated with the orthogonalized order flow innovation is interpreted as the proportion of FEVD explained by private information. The reported standard errors are analytic for the long-run price impact and Monte Carlo standard errors calculated using 1000 repetitions for the FEVDs.

The results suggest that a one standard deviation tradeflow or signflow shock results in a depreciation of the currency, since the long-run price impact is negative. However, it should be noted that none of the long-run price impacts are significant at the 1% level and only two are significant at the 5% level. The FEVDs are more conclusive, with customer signflow and tradeflow having a statistically significant double-digit information share (between 26 and 38%) at a monthly frequency. The FEVD at daily and weekly frequencies is less impressive, although it is

Table 5. Long-run impact of customer order flow.

Column (1) shows the long-run price impact of an orthogonalized 1 standard deviation shock to orders. For signflow a shock of +1 represents positive net cumulative order flow that day, week or month. For tradeflow a shock of +1 represents a single buy order. Column (2) shows the variance decomposition of the RETURN equation with the flow variable logically preceding the return variable in the Cholesky decomposition. Daily Vector Autoregressions (VARs) are run with five lags of each variable. Weekly VARs are run with four lags of each variable. Monthly VARs are run with two lags of each variable. Standard errors are shown below each estimate in parentheses. *, ** and *** show significance at the 10%, 5% and 1% levels, respectively.

	Flow variable	Long-run price impact (1)	Var decomp, flow first (2)
Daily	Signflow	−3.2606 (5.29006)	0.9954 (0.57212)
	Tradeflow	−18.23567** (7.49643)	5.871255*** (1.28952)
Weekly	Signflow	−44.53209* (29.1142)	4.395273* (2.63464)
	Tradeflow	−45.88572 (34.1141)	11.24987*** (3.6759)
Monthly	Signflow	−229.2034** (105.653)	26.12896** (9.59696)
	Tradeflow	45.064 (143.635)	38.42889*** (9.9612)

still statistically significant when the tradeflow variable is used. It appears, therefore, that while customer trading does not have a large price impact it is important for price discovery. These findings are consistent with Payne (2003) who finds that when liquidity is high, individual trades have small permanent effects but large information effects. Hence, customer trades may be important not because they move markets, but because they may hint at where the market is likely to go in the future.

The fact that customer trading is particularly informative at a monthly frequency appears to be consistent with Evans and Lyons (2005), who posit that customer order flow is informative at lower frequencies. Evans and Lyons (2005) argue that in the presence of dispersed information, order flow aggregated at lower frequencies gives a stronger signal of current and future changes in macro fundamentals than lagged macro variables.

We also estimate a two-equation VAR with customer data split by economic sector. Daily, weekly and monthly analysis is performed with both signflow and tradeflow variables. The defining feature of the customer data analysis is that only the diversified sector shows consistent evidence of a significant long-run price impact. The price impact ranges (depending on the flow variable used) from 11–17 tick depreciation in response to net buying at a daily frequency through to approximately 50 tick depreciation response to net buying at the weekly frequency. This is consistent with the regression results and the theory that diversified customers are liquidity-providing pull customers.

Other economic sectors do not appear to have any systematically significant price impacts. Given that the diversified sector makes up the majority of customer data (67% of all customer trades), it is interesting to note that the estimation results are much stronger when this sector is isolated. It may be that the financial sector, which dominates the data in terms of value of trades, is washing out the price impact in the aggregated VAR shown in Table 5.

The FEVDs show that diversified order flow has a high single-digit information share of the forecast error. However, the financial economic sector also appears to have a significant information share, particularly at the monthly level, which is consistent with Payne (2003).

The miscellaneous sector also shows consistently significant information shares across the estimated VARs. This could arise from the miscellaneous sector containing diversified firms.[6]

5. Summary

We test empirical implications of the Evans and Lyons (2002) three-period trading model. Specifically, this model predicts a long-run equilibrium relation between the exchange rate and customer order flow. Our results show no evidence of a cointegrating relation between these two variables. Since the two non-stationary variables (FX rate and cumulative order flow) are not cointegrated, we are not able to estimate an error correction model and therefore use OLS regression to estimate a relation between FX returns and net order flow.

The Evans and Lyons (2002) model posits that there are differing effects between financial and diversified order flow. In line with the predictions of the model, we find no significant relation between aggregate order flow and the exchange rate at any frequency. However, when we separate the financial and diversified order flow these results change. Specifically, the coefficient on diversified order flow is significantly negative, which is consistent with the model. This result suggests that such customers are liquidity providers and thus should be negatively correlated with the exchange rate. However, we do not find the complimentary finding of a significantly positive relationship between financial order flow and the exchange rate. The estimated coefficients tend to be positive but are not significant.

The limited support for the trading model is further investigated using VAR techniques. The analysis suggests that there is no significant dynamic relation between customer order flow and FX returns.

Overall, these results suggest that the non-financial customer order flow of a typical commercial bank does not carry much information about FX prices.

Acknowledgements

We are especially grateful for helpful comments from Igor Geninson, Tom Smith, Jian-Xin Wang and Minxian Yang.

Funding

This work was supported by ARC Linkage (grant no. LP0882468).

Notes

1. The large quantity of swap transactions merits further discussion. FX swap trades appear to be often used to implement a swap-forward or tom-next roll. This is a way for customers to borrow money for short periods of time by purchasing spot currency and immediately selling it forward for later delivery (usually 1 or 2 days forward). Since these yield zero net order flow they are usually removed from consideration.
2. The results for regressions where returns are computed at a weekly frequency are qualitatively similar (and therefore not shown). Total aggregate order flow shows no correlation with the exchange rate. The previous caveat about signflow and tradeflow variables being dominated by the diversified sector also applies. Financial firms' order flows are insignificantly positively related to the exchange rate, except when the tradeflow variable is used. Diversified and the similar miscellaneous flows are again significantly negatively correlated with the exchange rate using all trade variables. Finally, government flows are insignificantly correlated with the exchange rate variable, with a sign that varies depending on the trade variable used.
3. The UP1 and UP2 sub-period weekly results are consistent with those for the daily analysis, while during the DOWN sub-period, the weekly results differ from the daily analysis – the coefficient on the financial flow variable is insignificantly negative. If there is push–pull customer separation we would expect this coefficient to be positive.
4. A possible drawback of the data is that it represents a relatively small slice of the total AUD/USD market. Given this, one would expect that the customer order flow of any one bank is not likely to be completely informative, particularly given its unique composition of firms. Hence, it may be of interest to consider a different, smaller market with fewer key players. Relative to the AUD/USD market, the bank has a larger share of the smaller AUD/New Zealand dollar (NZD) market. We repeat the regression analysis for the AUD/NZD exchange rate and investigate differences between customer groups. The statistical properties of the AUD/NZD data are very similar to the AUD/USD data. It is therefore interesting to note that there is no significant relation between the mid-quote return and any trade flow variable from any sector at any frequency. The results for the weekly and monthly tests are not shown but are qualitatively similar. Hence it could be that price discovery in these currencies is taking place through AUD/USD and NZD/USD trade, rather than through AUD/NZD trade.
5. The number of lagged variables in the VAR model itself is not chosen according to an information criterion; rather it is based on economic intuition and data limitations. The results do not qualitatively change with the use of an information criterion-defined number of lags.
6. We also conduct the VAR analysis by sub-period. The sub-period results almost universally suggest insignificant price impacts and information shares and are not shown. Further, the sub-period VAR is strongly affected by data limitations (particularly the DOWN and UP2 sub-periods) at all but the daily frequency.

References

Bank for International Settlements (2007) *Triennial Central Bank Survey - Foreign Exchange and Derivatives Market Activity in 2007*. Basel, Switzerland: BIS Press & Communications.

Berger DW, Chaboud AP, Chernenko SV, Howorka E and Wright JH (2008) Order flow and exchange rate dynamics in electronic brokerage system data. *Journal of International Economics* 75(1): 93–109.

Bjønnes GH and Rime D (2005) Dealer behavior and trading systems in foreign exchange markets. *Journal of Financial Economics* 75: 571–605.

Bjønnes GH, Rime D and Solheim HOA (2005) Liquidity provision in the overnight foreign exchange market. *Journal of International Money and Finance* 24: 175–196.

Carpenter A and Wang J (2003) Sources of private information in FX trading. Working paper. Australian School of Business.

Cheung Y-W and Chinn MD (2001) Currency traders and exchange rate dynamics: A survey of the US market. *Journal of International Money and Finance* 20(4): 439–471.

Danielsson J and Love R (2006) Feedback trading. *International Journal of Finance and Economics* 11(1): 35–53.

Evans MDD and Lyons RK (2002) Order flow and exchange rate dynamics. *Journal of Political Economy* 110: 170–180.

Evans MDD and Lyons RK (2004) A new micro model of exchange rate dynamics. NBER Working Paper No. 10379.

Gereben A, Gyomai G and Kiss N (2006) Customer order flow, information and liquidity on the Hungarian foreign exchange market. Magyar Nemzeti Bank, MNB Working Paper.

Hasbrouck J (1991a) Measuring the information content of stock trades. *Journal of Finance* 46: 179–207.

Hasbrouck J (1991b) The summary informativeness of stock trades: An econometric analysis. *Review of Financial Studies* 4: 571–595.

Killeen W, Lyons R and Moore M (2006) Fixed versus flexible: Lessons from EMS order flow. *Journal of International Money and Finance* 25: 551–579.

Kyle AP (1985) Continuous auctions and insider trading. *Econometrica* 53: 1315–1335.

Marsh IW and O'Rourke C (2005) Customer order flow and exchange rate movements: Is there really information content? Cass Business School Research Paper, London.

Osler CL (2008) *Foreign exchange microstructure: A survey of the empirical literature*. Typescript, Brandeis University, Waltham, MA.

Payne R (2003) Informed trade in spot foreign exchange markets: An empirical investigation. *Journal of International Economics* 61: 307–329.

Reitz S, Schmidt M and Taylor M (2007) End-user order flow and exchange rate dynamics. *Deutsche Bundesbank Discussion Paper Series 1: Economic Studies.* Frankfurt am Main, Germany.

Treepongkaruna S and Gray S (2009) Information and volatility links in the foreign exchange market. *Accounting and Finance* 49: 385–405.

Treepongkaruna S, Brooks R and Gray S (2012) Do trading hours affect volatility links in the foreign exchange market? *Australian Journal of Management* 37: 7–27.