Individual trader records on EBS Spot: what do they tell us about private information in FX trading?

Michael J. Moore and Richard Payne*

Abstract

We analyse novel data from the dominant inter-dealer FX brokerage, containing dealer-level identifiers. We argue for an information interpretation of the link between order flows and exchange rates by demonstrating that the information content of trading varies with dealer trading style. In liquid dollar rates, aggressive deals by traders who specialize activity in a pair are most informative, while in cross-rates triangular arbitrageurs are best informed. In the latter, the ability to forecast flows explains the information advantage of triangular arbitrageurs. In liquid pairs, specialists can both forecast flows and the component of the return that is uncorrelated with flow.

*Moore, Queens University Belfast and Harvard Center for Population and Development Studies; Payne, Warwick Business School. Correspondence: mjmoore@hsph.harvard.edu, tel: (+)1-617-496-4573, fax: (+)1-617-495-5418. Thanks to ICAP-Electronic Broking Services for providing the data used in this study. Moore thanks the Economic and Social Research Council for financial assistance (grant number RES-000-22-2177). Many thanks to Dagfinn Rime, Carol Osler, Ian Marsh, Maurice Roche and seminar/conference participants at University of Madison-Wisconsin, the 2009 EMG Microstructure conference at Cass Business School, the ICMA Centre at the University of Reading, the Institute of Finance conference at University of Leicester, Union College, New York and the Federal Reserve Bank of New York for helpful comments and suggestions.

1 Introduction

The foreign exchange market is the world's largest by any standard. Daily turnover in April 2007 for spot, forward and swap transactions alone is reported by the Bank for International Settlements at \$3.21 trillion. An additional \$0.29 trillion was traded daily in over-the-counter foreign exchange derivatives (Bank for International Settlements 2007). The comparable figure for US treasuries is \$0.20 trillion while daily turnover on the NYSE seems puny at as little as \$0.05 trillion (see Osler (2008)). There should be little doubt that FX markets are greatly important, especially given the central role foreign exchange rates play in international macroeconomics.

Over the last decade or so, academic understanding of exchange rate determination and currency trading patterns have developed greatly due to a literature at the interface between microstructure finance and international macroeconomics. This literature, sometimes referred to as the "new micro approach" to exchange rates, focusses on the linkages between trading activity in FX markets and exchange rate changes.[1] Early examples of this work include Lyons (1995) and Lyons (1997). The core result from this literature is the following: a strong positive correlation exists between order flow, defined as the number (or aggregate quantity) of buyer-initiated less seller initiated trades in a given interval of time, and exchange rate changes measured over the same interval. While early demonstrations of this result were at the intra-day sampling frequencies common in microstructure analysis (see, for example,

[1]Dagfinn Rime maintains a bibliography of research related to information and order flow effects at `http://www.norges-bank.no/upload/import/english/research/dri/fxmicro-biblio.pdf`.

Payne (2003)), the breakthrough paper for the order flow approach was Evans and Lyons (2002), in which it was shown that order flow has explanatory power for spot exchange rate returns at sampling frequencies relevant to macroeconomics.

The finding that order flows have strong explanatory power for returns has proved to be remarkably resilient. It holds across sampling frequencies, across exchange rates and across different historical time spans.[2] However, there is, as yet, no consensus as to the source of this result. One branch of the literature asserts that order flows move exchange rates as *flows contain information*. However, in the FX context, it is not clear what this information might be. That it is 'inside' information about future 'fundamentals', such as monetary policy or macroeconomic announcements, seems to be, at first sight, unlikely. However, in a series of papers, Evans and Lyons, both separately and together, have developed a theoretical approach which emphasises the revelation of dispersed information about both discount rates and fundamentals through order flow. FX trading might reveal private information about idiosyncratic export demands, which are only aggregated and published in official statistics much later on. However, FX order flow may not just anticipate future common knowledge information. The 'portfolio shifts' approach argues that private information about changes in risk aversion and liquidity preference, for example, which will never become common knowledge, can only be impounded in price through the trading process itself.

[2]Further, it holds across asset classes. Research in equity and bond markets which follows the same path as the earlier work in FX finds similar results. See, for examples, Chordia, Roll, and Subrahmanyam (2001) and Vega and Pasquariello (2007).

An alternative view is that the relationship between returns and flows is due to the transitory impact of flows on liquidity and that flows are unrelated to information. Papers such as Breedon and Vitale (2004) and Bacchetta and van Wincoop (2006) argue along these lines. Berger, Chaboud, Chernenko, Howorka, and Wright (2008) aggregate order flow for EUR/USD and USD/JPY to monthly frequencies and find that the long term cointegrating relationship between exchange rate levels and cumulative order flow observed by Bjonnes and Rime (2005) and Killeen, Lyons, and Moore (2006) becomes weak or non-existent. This is interpreted as evidence that the effect on order flow on returns is due to a liquidity effect that is slow to dissipate. However, Chinn and Moore (2009) provide a robust response to this evidence, showing that, in monthly data, the cointegrating relationship is only apparent if the traditional monetary and real fundamentals are added to the cointegrating vector alongside cumulative order flow.

The contribution of this paper is to provide fresh evidence for the information interpretation of the relationship between foreign exchange order flow and exchange rate returns. To accomplish this, we make use of a unique new data set. Our data contains a complete record of all trades executed on the largest consolidated segment of the spot FX market, the electronic inter-dealer market run by EBS, for one month in 1999. Most importantly, we are given anonymized identifiers for the counter-parties to each trade. For each execution we have a record of the institution and the individual trader IDs on both aggressive and passive sides. This is the first study that uses dealer-level trading records for an entire trading platform at the microstructure level to measure differences in the information content of FX trades. Previous studies of similar kind in equity markets include Hau (2001) and Linnainmaa (2007). A recent related paper in the FX literature is Bjonnes, Osler, and Rime (2008) though they only examine the inter dealer records of a single bank. They stratify that bank's counter-parties by size and find that the larger the counter-party institution, the higher the price impact of its market orders. They go on to conclude that larger institutions tends to be better informed than their smaller counterparts.

We use our data on trader identities and deals to create a set of attributes for each trader in the data set. First, we measure the size of the dealing room in which each trader is located in order to discriminate between individuals working for large or small institutions. We proceed to classify traders according to the nature/style of their own activity. We identify, in liquid dollar pairs (EUR/USD, USD/JPY and USD/CHF), *specialist* traders who do the vast proportion of their trading in a single liquid rate. We then isolate traders who focus on arbitrage across a currency triangle. For example, a Yen *triangular* trader, is active in all three of the following rates: USD/JPY, EUR/JPY and EUR/USD. Finally, we separate out *high volume* or *big* traders (who have not been previously classed as specialists or triangular traders), leaving a tail of *other* traders who have not been placed into any of the preceding three trader-style categories.

We then study the how the price impact of trading varies with the classification of the aggressive and passive traders involved, and derive the following results. First of all, traders working on large floors move prices significantly further than those on small floors when they trade aggressively i.e. an aggressive buy (sell) trade by a trader on a large floor moves prices upwards (downwards) by more than an aggressive buy (sell) from a trader on a small floor. When a trader on a large floor executes passively, however, prices move significantly less far than if the passive counter-party is on a small floor. Thus, if one aggressively buys currency from a trader on a large floor, prices increase, but by a smaller amount than if the counter-party was on a smaller floor. These results suggest that traders on large floors are better informed than others – their aggressive trades move prices further and the price-moves against them after they trade passively are smaller in magnitude.

We further demonstrate that, in liquid rates, specialist traders appear to have superior information to other traders. They move prices significantly further than any other class of trader when trading aggressively and suffer smaller adverse price movements when trading passively. In less liquid cross-rates, the price impacts provide evidence that the traders with the greatest informational advantage are the triangular traders. It should be noted that the triangular traders' advantage is limited to the cross-rate only and does not extend to the liquid legs of the triangle.

Finally, we show that high volume traders who have not fallen into either of the specialist or triangular trader categories show an information advantage over the class of 'other' traders.

We interpret the results above as supporting the information-approach to order flow. Our reasoning for this is as follows. There is, a priori, no economic reason to think that the price move following a trade should vary across trader types for liquidity reasons. If a trade conveys no information (about future fundamentals, discount factors or future market activity), then that trade should move prices by the same amount regardless of the identities of the counter-parties — a \$1mn aggressive buy from an individual who trades once a day and sits on a small floor should have the same liquidity effect as a \$1mn buy from a dealer on a huge floor and who trades once every minute. Thus, evidence that price impacts vary systematically across trader types is supportive of the information content of order flow. Effectively, we provide evidence that certain trader classes are better able to forecast, and then trade upon, short-run price movements than others and equate this variation in forecasting power with variation in information.

This argument above naturally leads one to inquire as to the nature of the superior information that specialist traders or traders on large floors possess. We evaluate one possibility in the current analysis, that better informed traders are better able to forecast *order flow*. In a world where order flow and price changes are contemporaneously correlated, a trader who can predict and then trade ahead of flows should be able to derive a positive return from this activity. We show that, in most cases, traders who have been previously identified as better informed can indeed better forecast flows, but that this ability does not fully explain the returns around their trading activity. This result is most clear for the most liquid rates. Thus, in the liquid rates, superior information is something more than superior market intelligence regarding future flows.

The plan of the paper is as follows. The next section delves further into the institutional background to the data set that we employ. In section 3, the data is introduced and described. Section 4 analyses the data and presents the main results of the paper. The final section offers some concluding remarks.

2 Institutional background

The foreign exchange market is completely decentralized. There is no organized physical location at which trading takes place and the market is almost totally unregulated. It has a multiple dealership structure with extremely low transparency. Nevertheless, for one month every three years, the Bank for International Settlements carries out a census of foreign exchange activity in the major countries of the world.[3] From this, one obtains a good overview of the market. Three broad categories of trading, based on counter-party type, are identified. Inter-dealer trading amounts to 43% of the market, trading involving non-dealing financial institutions accounts for 40% while the remainder, 17%, comes from non-financial customers. The ultimate customer share has remained fairly stable over the years but the inter-dealer share has fallen from a peak of over 60% in 1998. The market segment that has grown to take up the slack has been the non-dealing financial institution sector. However the distinction between dealers and financial institutions is no longer clear-cut in FX. Some electronic brokerages, for example, now allow hedge funds and parts of the non-dealer FX professional trading community the access to their platforms such

[3]'Major' is very thoroughly interpreted: there were 54 countries included in the most recent survey in April 2007.

that many financial institutions have assumed the role of dealers. Our focus is unambiguously on the inter-dealer and financial institution part of the market, however it is divided. The customer segment is surveyed, inter alia, by Osler (2008).

The inter-dealer market trades over the counter (OTC), through direct inter-dealer trading, as well as via voice and electronic broking. However electronic communication networks of various kinds now dominate the market: see Barker (2007). There are two main electronic order-driven inter-dealer systems. The more important of the two is Electronic Broking Systems (EBS), now part of ICAP plc. Its specialities are the five pairs: USD/JPY, EUR/JPY, USD/CHF, EUR/CHF along with the anchor pair EUR/USD. Reuters Dealing 3000 is the primary liquidity source for USD/GBP, EUR/GBP, USD/AUD, USD/CAD and many lesser currency pairs.[4] Since our data come exclusively from EBS, we concentrate on a number of its characteristic features at this point.

EBS operates as an electronic limit order book with liquidity supplied via limit orders and liquidity demand via market order (and direct limit order crosses). The EBS platform only permits trading between counter-parties where there is sufficient bilateral credit. It displays to all its traders the 'EBS Best' price pair (i.e. the highest bid and lowest offer in the market at the time). Additionally each trader observes an idiosyncratic 'Best Dealable' price pair which comprises the highest bid and lowest offer that the dealer has access to, given his credit relationships with other institutions.

[4] At the inter-dealer level, most pairs are not traded at all. Each country typically has just one liquid traded pair, usually versus the dollar and when not, the euro. On the two inter-dealer systems, only a subset of these (approximately 30) offer serious liquidity. See LondonFX (2008).

Banks can only execute against 'Best Dealable' prices and, clearly, the 'Best Dealable' prices must be weakly inferior to the 'EBS Best' prices. For further details, see Ito and Hashimoto (2006). Pre-trade, the quantities available at both 'Best' and 'Best Dealable' prices are also visible on the screen to each trader. However the identities of the liquidity suppliers, even for the 'Best Dealable' quotes, remain anonymous.

Post-trade, both sides see each other's bank code and individual trader identity. However post-trade transparency for those not participating in a particular trade is extremely limited. EBS posts to the platform the last trade price (by currency pair) in each half-second time slice (if there is a deal) and whether that deal was an aggressive buy or sell. There is no other information offered: nothing about size nor any other deals in the time slice. A participating dealer is limited to extracting information from other screen information, for example attempting to infer size from variation in the quantities available for trade before and after execution.

3 Data

3.1 Data structure

The raw data used in this study consists of a record of all deals completed on Electronic Broking Services (EBS) for the calendar month of August 1999.

The foreign exchange trading day conventionally runs for the 24-hour period from 9pm to 9pm GMT. Trading is very rare between 9pm on Friday evening and the same time on Sunday evening and there are no such deals in our sample. Our data set begins at 9pm on Sunday 1st August and concludes at 9pm on Tuesday 31st August: allowing for weekends, this amounts to 22 days of trading. During our sample period, liquidity on EBS concentrated on five currency pairs. In four of these, USD/JPY, EUR/JPY, USD/CHF and EUR/CHF, EBS was almost the sole source of electronic inter-dealing trading, while it was the major source (in excess of 80%) of electronic inter-dealer trading in EUR/USD.[5] There are other currency pairs in the record but these are not analyzed because EBS was not, at the time, a major source of liquidity for them.[6]

Each record in our data set represents a deal with a timestamp, currency pair, direction of transaction, price, volume and deal counter-party information. To illustrate this, Table 1 reproduces the data for the six deals which took place on Monday 2nd August 1999 during a 10 second interval.

The first field in each line of data is the trade date which, in Table 1, shows as 08/02/99 – 2nd August 1999 – in each case. The second field is a GMT time stamp, expressed as 'hour:minute:second'. Thus the first deal in Table 1 took place at twenty seconds after one minute after midnight. The next deal took place four seconds later. There are two deals recorded for twenty nine seconds past the minute. Within EBS's confidential record there is a more refined millisecond time stamp to which we do

[5]Breedon and Vitale (2004), Table 3, report data on dollar/euro trades on both EBS and Reuters Dealing 3000 for the period August 2000 to Mid-July 2001. From the information that they report, we can calculate that EBS had a share of the electronically brokered inter-dealer market of 87.7% by value. Anecdotal evidence suggests that this fraction has been rising secularly.

[6]Recently, EBS has developed as a significant source of liquidity in pairs in which Reuters Dealing 3000 used to have a near monopoly. These include GBP/USD, AUD/USD and USD/CAD.

not have access. However, the true chronological ordering of deals is preserved in the data provided to us such that we can conclude that the USD/JPY deal was executed before the EUR/JPY deal. The next item is the currency pair. In the six deals in Table 1, there are three different pairs: EUR/USD, USD/JPY and EUR/JPY. The first currency mnemonic is the base currency while the mnemonic after the forward slash is the currency in which the price is expressed. So EUR/USD gives a Dollar price for trading a quantity of Euros.

The next part of the record gives information on the direction of trade: 'B' denotes a buyer initiated trade and 'S' a seller initiated trade. The price, given to five significant digits follows next. It is worth mentioning that the number of significant digits is not like a tick size constraint in equity markets. There is really nothing to stop FX traders from quoting as many digits as they wish. For example, since 2005, Barclay's Capital offers an additional significant figure in trading major currency pairs. However EBS sees no commercial need to do this.[7] The next field gives the traded volume in millions of base currency units. Trades on EBS are restricted to integer multiples of millions of currency units.

The last four fields provide the distinctive feature of the data set. They provide anonymous mnemonics for the identity of the traders. The first two 'Maker Bank' and 'Maker Trader' refer to the passive side of the transaction – the liquidity supplier. The two remaining fields 'Taker Bank' and 'Taker Trader' identify the aggressor, the dealer whose incoming buy or sell instruction crosses with a pre-existing order and

[7]For a discussion of this see Goodhart, Love, Payne, and Rime (2002).

thus triggers a trade. It should be noted that the 'Bank' identifiers refer not to a financial institution, but to a specific trading floor of a financial institution. Thus, the same financial institution would have different 'Bank' identifiers for its trading floors in different geographical locations.

3.2 Basic statistical analysis of returns and trading patterns

Figure 1 shows the evolution of each of the five exchange rates over our sample period. It can be seen from the Figure that this is relatively stable period for all pairs with no particularly abrupt movements in any of them. The largest cumulative returns, of around -5% across August 1999, were in EUR/JPY and USD/JPY.

Table 2, Panel A provides summary statistics for the returns on each currency pair in event time. We show the first four moments of returns as well as the first-order autocorrelation coefficient. There is evidence of excess kurtosis in all rates as well as negative autocorrelation (likely due to bid-ask bounce). Panel B shows the same statistics in calendar time, where we have chosen a 5 minute sampling frequency so that the exercise is meaningful for the less liquid pairs. The main difference between the results presented in Panels A and B is that the return autocorrelation disappears in calendar time.

Table 3 gives the numbers and cash value of trades for each rate, both in total and broken down into buyer and seller initiated components. It also shows the related statistic of the average time between trades, measured in seconds. The most liquid pairs are EUR/USD, USD/JPY and USD/CHF in that order. It should be noted, though, that USD/CHF is much less liquid than USD/JPY. As for the cross-rates, activity in EUR/JPY is significantly above that in EUR/CHF.

Figure 2 refines our presentation of trade frequency by examining the distribution of exchange rate activity by time of day. The least liquid time of day is between 9pm and midnight GMT: this is not surprising as it is roughly the time between the New York close and the opening of Asian markets. EUR/USD is fairly liquid throughout the 24-hour day: it has peaks during morning European trading and again during the overlap between the US morning and European afternoon. EUR/CHF and USD/CHF show very similar diurnal variation. USD/JPY and EUR/JPY are also liquid throughout the day but with three peaks: corresponding to the Asian, European and US mornings. Our observations are consistent with the evidence presented in Ito and Hashimoto (2006) on EUR/USD and USD/JPY.

Table 4 shows the distribution of trade size by currency pair. What is surprising here is the remarkable similarity of the distributions irrespective of currency pair, mindful of the fact that the average rate for EUR/USD for August 1999 was $1.06 per Euro. The majority of trades are sized at 1 to 3 million currency units. It is only in the case of EUR/CHF that we see any real deviation from the 'average' pattern. For this rate, there is evidence that trades tend to be somewhat larger than in other pairs, with around 13% of all trades being for at least EUR 5mn.

Our trade size results accord well with those from previous work. Hau, Killeen, and Moore (2002) have already shown that the mean trade size for EUR/USD was approximately $2 million dollars in 1999 and slightly higher for DEM/USD in 1998. The estimates in Table 2 of Bjonnes and Rime (2005) of trade size for electronically brokered deals for DEM/USD for 1998 are consistent with this. They also show that the mean trade size for the cross-pair DEM/NOK in March 1998 was about DM 3.8-4.1 million or also again just over $2 million per trade. Our results show that what Bjonnes and Rime (2005) implied about the distribution of trade size is quite general across currency pairs.

3.3 Banker and Trader Identities

The exciting feature of the data illustrated in Table 1 is the availability of bank and trader identities for both sides of every deal. It has already been emphasized in section 3.1 that the bank identity corresponds to a dealing room rather than a financial institution, per se. To be clear, the EBS bank code identifies all activity emanating from a specific financial institution and in a specific physical location. Thus, Bank X's Tokyo and London activity would fall under different bank identifiers. The trader identifiers then isolate individuals or desks within that location.

The data contains executions from 727 dealing rooms and 2867 traders and thus the average dealing room contains around 4 traders. The histogram in Figure 3 shows the distribution of traders by dealing room. Descriptive statistics for this distribution are given in Table 5. The distribution is markedly right skewed with the bulk of dealing rooms being quite small: the number of rooms with only one identified trader is startling.

The immediate issue to raise is how important are the small relative to the large dealing rooms? Panel A of Table 6 shows the distribution of trades, separately by currency pair, across the four quartiles of dealing room size. Not surprisingly, larger dealing rooms have a greater share of overall trading activity. Panel B of Table 6 addresses a more precise question. It shows the distribution of trades, per trader across dealing room size quartiles. The table suggests that most of the casual traders are employed in modest dealing rooms. Finally, Panel C of Table 6 shows the distribution of trade size in each currency pair across dealing rooms. It is clear that the largest trades are carried out from the biggest dealing rooms. Overall, Table 6 shows that the most intensive traders work on large trading floors.

This analysis of trader activity broken across different sized dealing rooms leads naturally to a study of which traders are most active overall. To this end, we examine the frequency with which a trader executes deals. Table 7 breaks traders into five categories. The most active traders are those that trade at least 5 minutely, then those that trade up to fifteen minutely, up to hourly, daily and less frequently than daily. It is somewhat arbitrary to label a trader as 'big' or 'small' but the authors' experience of FX markets suggests that it is hard to characterize an FX trader as 'big' if (s)he is not trading at least once every quarter of an hour. Indeed, this may be too generous but for the rest of this paper, we describe traders in the first two rows of Table 7 as 'big' traders. This category is significant because we hypothesize that the capacity to observe a significant chunk of order flow is a major source of information for traders.

While trading frequency might well be an important proxy for information quality, one might also regard 'specialism' of research and trading in a specific security as being important for generating an informational advantage. To this end Table 8 classifies the subset of 'big' traders (defined according to their total activity across all 5 rates) by the proportion of their own activity that is executed in each exchange rate. Via this table, we label trader Z as a specialist in a given exchange rate if at least 90% of all of that trader's deals are in the specified rate.[8] Table 8 demonstrates a clear mass of traders in EUR/USD and USD/JPY who execute more than 90% of their trades in those rates and thus can be justifiably labelled as specialists. Such a clear set of specialists is harder to identify for the USD/CHF, with only 5 traders doing over nine tenths of their activity in that rate and there being no obvious mass of probability at relatively high specialization levels. Thus it could be argued that, particularly for USD/CHF, a less stringent specialist definition might be justified. However, for consistency, we use the 90% criterion for all three liquid dollar pairs.

For the two cross-pairs, it is clear from Table 8 that there are virtually no 'specialists' in the sense that we have just defined. As such, we ignore this classification for the crosses. Instead, for the cross rates we look to triangular exchange rate arbitrage to sub-classify the 'big' traders. We argue that a trader in the cross-rate should not focus only on direct trade in the cross-rate, but also on indirect trade between the

[8]We think that specialist is an accurate term to use to describe these types of traders. We are aware that this term is also the name of designated NYSE market makers. Such individuals provide both dealership and brokerage services and, in return for these services, the exchange grants the specialist a right to make a market in a particular stock. For a full discussion, see Benveniste, Marcus, and Wilhelm (1992). There is, of course, no analogy being made between the two types of trader.

same two currencies that is accomplished using the two liquid legs of the currency triangle. For example, a trader in EUR/JPY must be aware of direct EUR/JPY flows *and* must also keep an eye on flows in EUR/USD and USD/JPY as they give an indirect, but very liquid, route for trading Euros for Yen.

Recent work by Lyons and Moore (2009) provide a theoretical analysis of the exploitation of private information in currency triangles which shows how all three legs of a currency triangle are used in equilibrium and our empirical approach builds on that work. We define a 'triangular' trader as one who is active in all three legs of a triangle. Our definition requires such a trader to conduct at least 15% of her trading activity in each of the three legs of a currency triangle.[9] In terms of Table 8, this means that we exclude traders in the bottom one and top three deciles altogether. The JPY triangular traders are found at the intersection of the deciles between 10% and 70% in the three columns headed EUR/USD, USD/JPY and EUR/JPY and equivalently for the CHF triangle. It is not possible to read off these numbers directly from the table but the numbers of JPY and CHF triangular traders using this criterion are 31 and 29 respectively.

Thus four categories of trader arise out of our classification scheme;

- 'Specialist' (S) traders in EUR/USD, USD/JPY and USD/CHF who trade frequently and focus their activity largely in one rate.

- 'Triangular' (T) traders who are arbitrage traders in all three legs of a currency

[9]As with other definitions, our results are robust to reasonable variations in the activity cutoff used in this definition.

triangle and who also trade frequently.

- 'Big' (B) traders are frequent traders who have neither been classified as 'Specialist' nor 'Triangular'
- 'Other' (O) traders, who execute infrequently.

Table 9 summarizes the classification of traders, providing the number of traders in each category by trading pair. Trivially, the largest category is 'Other' because most traders are small. Note that the number of Big traders is a significant proportion of total trader numbers, as is the number of Specialists in EUR/USD and USD/JPY.

In the tables that follow we have abbreviated the names of our four trader categories with their initial letters. Thus we have Big (B), Specialist (S), Triangular (T) and Other (O) traders. We also classify each deal according to the category of the Maker (M) and the Taker (T). Thus, we will often present statistics separately for the eight combinations of trader category and maker/taker distinction, with the eight being BM, SM, TM, OM, BT, ST, TT and OT. To be explicit, a trade labelled as ST has a 'Specialist' trader on the taker side and a trade denoted OM has, on the maker side, one of the set of 'Other' traders.

4 Analysis and Results

Our core hypotheses relate to variation in access to private information across our various trader categories. We hypothesize that the scale and concentration of a trader's own order flow will positively affect his or her information quality. Similarly, a trader's environment, specifically the size of the trading floor on which that trader is located, will also affect information quality — a trader on a large floor should have access to better information than a trader on a small floor. This leads us to the following empirical hypotheses;

- In all currency pairs, the price impact associated with trading should be larger when the aggressive counter-party is on a large floor than it is when the aggressive trader is on a small floor. Conversely, when the passive counter-party is on a large floor, we would expect price impact to be smaller than when the passive counter-party is on a small floor.

- In liquid pairs we expect the trades of aggressive Specialist dealers to exert larger effects on prices than trades of other dealer categories, and the passive trades of Specialists should be associated with the smallest mean price impacts of all categories.

- In the cross-rates we expect aggressive (passive) trades by Triangular traders to have the largest (smallest) price impacts of the three trader classes.

- In all rates, we expect aggressive trades of Other dealers to have the smallest price impacts and the passive trades of Other traders to have relatively large impact when compared to the passive trades of the rest of the trader classes.

4.1 Trade size

It is well known from equity markets that informed traders manage trade size as a mechanism to conceal information. For example, Chakravarty (2001) analyses how informed equity traders concentrate their activity in medium sized trades. In the foreign exchange market, Bjonnes and Rime (2005) show that trade size is related to the information content of spreads. It seems natural, therefore to examine differences in mean trade size across our eight combinations of trader category and market/taker classification.

Table 10 displays information on trade size by trader class. For the liquid Dollar rates, it is clear that specialist and big traders trade in larger quantities, whether actively or passively. For the cross-rates the big traders are clearly those who deal in greater size. While, as the test statistics contained in the table indicate, the differences in mean quantities dealt across trader types are statistically significant, they are economically relatively small. These results suggest that there is little ability to discriminate between trader classes based on deal size and thus in the following analysis of the price impact of trades we abstract from trade size altogether.

4.2 Price Impact of Order Flow

4.2.1 Empirical methodology

Our method for investigating the information content of trades departs from the method used elsewhere in this literature (e.g. Bjonnes, Osler, and Rime (2008)), where exchange rate returns over regular calendar time intervals are regressed on aggregate order flows over the same intervals. Rather than aggregating deals, we examine the price movements surrounding individual trades. As such, our methodology is closer to that in the empirical microstructure end of the order-flow literature (e.g. Payne (2003)).

To begin, for each deal in the sample we define the following price impact measure;

$$\Delta p_{c,i,h,h'} = 10000 \times d_{c,i} \times \ln\left(\frac{p_{c,i+h}}{p_{c,i-h'}}\right) \tag{1}$$

where the subscript c indexes the currency pair, the subscript i indicates that we are looking at the ith deal in that currency pair and h and h' are positive integer parameters. The variable $d_{c,i}$ is a trade direction indicator which is +1 for an aggressive buy trade and -1 for an aggressive sell. Thus, our price impact variable is just the change in price, measured in basis points, from h' trades before to h trades after deal i in currency pair c. The trade direction indicator ensures that both buy and sell trades should have positive mean price impact. A similar construction is used in Linnainmaa (2007).

In what follows we set the parameter h' to 5 and allow h to vary between 5 and 50 trades. Order flow in our data is positively autocorrelated because of what Biais, Hillion, and Spatt (1995) call the diagonal effect. By choosing h' to be larger than one, the effect of autocorrelated order flow on price impacts is filtered out. Setting h substantially above zero allows for delayed adjustment of prices to the information contained in the current trade.

Our price impact variables are then employed as the dependent variable in the regression specification shown in equation (2). This equation conditions price impact on three sets of factors. First, impact is allowed to vary with the types of counter-party on the maker and taker side of the trade. Second, price impact depends on the size of the trading rooms in which both maker and taker sit. Finally, price impact varies with the evolution of the market-wide information environment.

$$\Delta p_{c,i,h,h'} = \sum_k \sum_j \beta_{j,k} X_{i,j,k} + \lambda_M MROOM_i + \lambda_T TROOM_i + \gamma \sigma_i + \delta \sqrt{DUR_i} + \epsilon_{c,i,h,h'} \tag{2}$$

In equation (2), variation of impact with trader type is permitted via the $X_{i,j,k}$ regressors. $X_{i,j,k}$ is a dummy variable which takes the value unity if and only if the aggressive counter-party of deal i was of type j and the passive counter-party was of type k. In liquid rates, both j and k can be one of Big, Specialist, Triangular and Other such that for these currency pairs there are 16 individual dummy variables. In the cross-rates (EUR/JPY and EUR/CHF), j and k can only be Big, Triangular or Other, such that there are 9 trader-type dummies in these regressions. In all regressions and for all combinations of trader types, we expect the price impacts from trading to be positive such that we expect $\beta_{j,k} > 0$ for all j and k. Our preceding hypotheses regarding the differential information quality of various trader types give a set of size inequalities that we would expect the $\beta_{j,k}$ coefficients to obey.

The variable $MROOM_i$ is the square root of the size of the room in which the maker of trade i is located and $TROOM_i$ is the square root of the size of the taker's room location.[10] We expect room size to have a positive effect on information quality and/or quantity, such that if a trade's taker (maker) is located in a large room, price impact should be relatively high (low). Thus, we expect $\lambda_T > 0$ and $\lambda_M < 0$.

Finally, we wish to control for time variation in the fraction of informed trades in the market, as the higher the common knowledge probability of an informed trade, the higher the price impact for all trade types. Linnainmaa (2007) accomplishes this using the bid-ask spread, which is increasing in the probability of an informed trade. In the absence of market spread data, we rely on the work of Bollen, Smith, and Whaley (2004) who argue that that the combined adverse selection and inventory components of the spread can be approximated by a variable which is directly proportional to $\sigma\sqrt{t}$, where σ is the standard deviation of the spot return and t is the time between trades. In equation (2), volatility and duration enter separately. This allows us to account for the results contained in Dufour and Engle (2000) who argue that the price impact of trading is decreasing in the time between trades, a result which goes in the opposite direction to the spread effect. Consequently, though we anticipate $\gamma > 0$, the sign of δ is indeterminate, a priori. Empirically, we measure volatility as the realized volatility in the hour preceding the trade based on minutely sampling. $\sqrt{DUR_i}$ is measured as the mean of the square root of all durations in the exchange rate under study in the hour prior to the current trade.[11]

[10]We use the square root of the number of traders as our measure of room size to allow for diminishing returns to scale. Both variables are demeaned before inclusion in equation (2).

[11]Both volatility and duration regressors have also been demeaned prior to inclusion in equation

4.2.2 Price impact results

The results of estimating these equations for a tick-time horizon of $h = 10$ trades are presented in Table 11.

Focussing first on the coefficients on those right-hand side variables not based on trader identity information, one can see that the volatility and duration regressors are positive and significant in every case. Thus trades tend to have greater impact in more volatile times, presumably due to this volatility representing increased information and inventory risk, and also tend to have larger impact when durations are high. This latter result runs counter to that presented in Dufour and Engle (2000).

Looking next at the coefficients on maker and taker room size, again our priors are confirmed. All coefficients are correctly signed and strongly statistically significant. Aggressive traders in large rooms tend to move prices further than aggressive traders in small rooms. Conversely, the deals of passive traders in large rooms tend to generate smaller price impacts than those of passive traders in small rooms. Thus, we obtain clear evidence that the immediate environment of a trader is important in determining her effect on prices. Larger trading rooms are likely to be associated with concentration of market knowledge and research skills, and possibly confer access to larger, or higher quality, customer flows. All of these these confer informational advantages to the traders located therein.

The coefficients which are most relevant to our hypotheses, however, are the price impact estimates for the various trader type indicator variables. These are measured

(2).

in basis points per trade. Across all rates, all of these coefficients are positive and all but one statistically significant (most very strongly so). At first sight, it may seem hard to identify a pattern among the sixty six estimated coefficients. However, there is a very strong pattern that is supportive of the hypothesis that some agents are systematically better informed and, more precisely, that this information is extracted from the trading process itself.

In Table 12, the estimates of Table 11 are averaged in a simple manner to reveal this pattern. Consider the column in Table 12 labelled EUR/USD. The entry of 0.27 for BM is obtained as the unweighted mean of the entries for the four rows of Table 11 in the EUR/USD column that involve a Big Maker (i.e. those labelled BMBT, BMST, BMTT, BMOT). It thus represents the 'average' price impact of trades in EUR/USD where the liquidity supplier is Big. An analogous interpretation can be given to the entries for SM (Specialist Maker), TM (Triangular Maker) and OM (Other Maker). From the perspective of the maker of the deal, a large, positive value for this price impact is unwelcome as it indicates that the average response to a purchase (sale) is a large drop (rise) in price.

Looking first at EUR/USD, trades of Specialist and Big Makers are associated with the smallest price impacts (in that order) with the mean impacts of trades of Triangular and Other Makers being much larger. A similar pattern exists for USD/JPY. In the case of USD/CHF, trades of Triangular liquidity suppliers have relatively low associated impact but other than that the pattern identified previously prevails in all liquid pairs.

Staying with the three liquid pairs, let us turn to the final four rows of the table, which look at the aggressor classifications; BT (Big Taker), ST (Specialist Taker), TT (Triangular Taker) and OT (Other Taker). Here, a large positive coefficient indicates a large rise (fall) in price subsequent to a purchase (sale). In the EUR/USD column, it is clear that the Specialist has the biggest impact followed by the Big traders with the Triangular and Other traders (in that order) bringing up the rear. The same ordering applies for aggressive deals in USD/JPY and USD/CHF.

Thus, for the liquid rates, the orderings across making and taking are entirely consistent. When supplying liquidity (i.e. when acting as 'makers'), deals by Specialist traders are associated with the smallest movements in price, but when Specialists take liquidity their trades clearly move prices further than do the trades of others. It would thus seem reasonable to assert that in the liquid pairs, Specialist activity carries most information, followed by Big, Triangular and Other traders in that order.

To add some statistical weight to this assertion, Table 12 also contains test statistics for the null hypothesis that the mean impact for a trade made by a Specialist is identical to that of the Big and Other trader classes. Similar statistics are presented for takers. The differences between impacts of Specialists and the other pair of trader types are of the expected sign in all but one case and are significant in 9 of 12 cases. The test statistics are most significant for the most liquid rates (EUR/USD and USD/JPY) and are stronger when comparing the taker side of trades than when comparing makers. Thus, perhaps unsurprisingly if they are trading on information, Specialists' aggressive price impacts are systematically stronger than those of the other two types of trader. Thus, not only are Specialists different from the class of Other traders, but their deals have significantly different price effects than do those of Big traders.

It is similarly easy to identify the winners and losers in trading in the non-dollar cross pairs. The Triangular trader type is now associated with the highest aggressive impacts and the smallest impacts when making. Aggressive deals from the Other trader class move prices by the smallest amount and their passive trades are associated with the largest price impacts. Table 12 again shows that, in the cross rates, impacts of Triangular traders are significantly different from those of both Big and Other traders.

In sum, our results suggest that, in liquid exchange rates, Specialist traders are those whose trades best exploit short-run trends in price. Their aggressive trades anticipate relatively large movements in prices, in that they buy (sell) before price increases (declines). Their passive trades are timed so that, subsequent price moves *against* them are limited. In the cross-rates, the Triangular traders take on the role played by the Specialists in the liquid rates. Across all rates, the class of Other traders displays the smallest ability to forecast short-run price changes. These results, strongly suggest informational differences between our trader categories in the sense that information is associated with forecasting power over future changes in price.

4.2.3 Sensitivity analysis

Our results are not peculiar to the event horizon that we have chosen nor to the fact that the regressions are estimated in event time. In Table 13, we report aggregated results from regressions estimated in calendar time, analogous to those in Table 12. Here the price impact variable is measured from 10 seconds before the trade ($h' = 10$) to 60 seconds ($h = 60$) after. The pattern in coefficients is remarkably similar to those from the tick-time analysis.

For the liquid dollar pairs, aggressive trades by specialists tend move prices furthest, followed by Big traders, Triangular traders and Other traders, respectively. This ranking is reversed when we look at passive trades. These differences are significant in 5 of 9 cases and more clearly so for the comparison of the aggressive trade impacts. For cross-rates, deals where triangular traders are taking have greatest impact, followed by Big and Other traders. Again, the ranking is reversed for the maker side of the market and all differences but one are statistically significant.

As a final robustness check, Table 14 reports mean tick-time impacts by trader type but where we have varied the post-trade impact horizon (h in equation (1)). In particular, we report results for $h = 5$ trades and $h = 50$ trades. Our basic results are, in broad terms, confirmed once more. Parameter estimates and the differences between trader classes are much more significant and impact rankings very consistent at the lower post-trade horizon ($h = 5$) but also persist to the 50 trade horizon. At the longer horizon, the results are less strong and much less significant for the less liquid exchange rates. This is perhaps unsurprising, as a 50 trade horizon covers, on average, between 20, for USD/CHF and 50, for EUR/CHF, minutes (see Table 3).

In sum, these checks indicate that our conclusions are not specific to the basic time-horizon over which we measure impact or to whether we conduct our analysis in clock or transaction time. Moreover, the broad consistency between the tick-time results based on 5 and 50 trade horizons strongly suggests that the impacts we are uncovering are information-based, and not due to transitory liquidity effects.

4.2.4 Order flow impacts by trader type

One potential explanation for our price impact results is that certain trader classes are better able to forecast order flow than others. Given that we know that there is a strong positive contemporaneous correlation between flows and returns, a dealer who can forecast flows should be able to position herself so as to profit from the price moves generated by that flow (buying when she anticipates positive flow and selling on negative flow forecasts).

To test this hypothesis, we run regressions similar to equation (2), but where the dependent variable is not the price impact of a trade, but the cumulative order flow surrounding a trade. Cumulative order flow ($F_{c,i,h,h'}$) is defined as the sum of signed traded quantity from h' trades prior to the current trade to h trades after. From this measure we remove the flow contribution of the current trade and we also delete any aggressive trades by either the maker or taker of the current trade. This filtering ensures that any systematic effects that we pick up are not due to dealers' ability to forecast their own flows.

$$F_{c,i,h,h'} = \sum_k \sum_j \beta_{j,k} X_{i,j,k} + \epsilon_{c,i,h,h'} \tag{3}$$

where the $X_{i,j,k}$ variables are as previously defined. We refer to the dependent variable in this regression at the tick-time flow impact. It is measured from 5 trades prior to the current trade to 10 trades after and is in millions of base currency units of the relevant exchange rate. Coefficient estimates of equation (3), aggregated across trader types as in the price impact regressions, are presented in Table 15. What this table makes clear is that those traders identified previously as dealing on information (Specialists in the liquid rates and Triangular traders in the cross-rates) have a significant advantage when it comes to forecasting order flows. It is also clear that such advantages are greatest and most significant in the less liquid cross-rates and for taking rather than making trades.

These results suggest that it is possible that the informational differences between trader classes previously identified is due entirely to a difference in their ability to forecast future trading pressures. Specialists in EUR/USD, for example, may outperform Other traders in terms of trading profits only because they have a better developed view of the order flows likely to arrive at market in the next few minutes or hours. This assertion is tested below.

4.2.5 Price impacts accounting for order flow

The results of the preceding subsection raise the possibility that the differences in price impacts across trader classes might be due to their differential ability to forecast order flows. We test this assertion by adjusting our regression specification to explicitly account for order flow on the right-hand side. Specifically, for each exchange rate and for $h' = 5$ and $h = 10$ we estimate the following regression;

$$\Delta p_{c,i,h,h'} = \sum_k \sum_j \beta_{j,k} X_{i,j,k} + \lambda_M MROOM_i + \lambda_T TROOM_i + \gamma \sigma_i + \delta \sqrt{DUR_i} + \kappa OF_{c,i,h,h'} + \epsilon_{c,i,h,h'} \tag{4}$$

where $OF_{c,i,h,h'}$ is cumulative order flow from h' trades before to h trades after the ith deal. This measure differs from that employed in equation (3) as trades made and taken by the counter-parties to the ith deal are *not* excluded from this flow computation. The change in the regression specification means that the estimated coefficients on the dummy variables ($X_{i,j,k}$) should now be interpreted as the effects of counter-party type on the price movement around the current deal, having already accounted for the effect of order flows.

We do not report individual estimated values of the coefficients from these regressions but, unsurprisingly, κ is positive and greatly significant in every case and drives all regression R^2 well over 0.10. The estimated effects of trade type on impacts are presented, in the usual aggregated format, in Table 16.

Table 16 makes clear that, for the liquid dollar rates, those traders previously identified as having the greatest informational advantages still do so after having adjusted for their ability to forecast order flows. Thus, Specialists do not just front run high-frequency order flows, they have some ability to forecast that part of the price move that is uncorrelated with flows. Again, this result is more pronounced on the taker side of trading. For the cross-rates, the informational differences identified previously disappear (in terms of statistical significance) when we account for the effect of flow on prices. For these less liquid pairs, our results suggest that the information advantage of the Triangular traders is generated entirely by their ability to forecast order flows.

Thus, for the most liquid exchange rates, the best informed trader class, the Specialists, displays two advantages: first they can better forecast future order flows and second they can predict that part of the high-frequency price change that is uncorrelated to order flow. For cross rates, our Triangular traders are best informed, but their advantage seems to rely solely on the ability to forecast future flows.

One might well wonder about the nature of the advantage which enables Specialists and Big traders in liquid rates to predict returns independently of their ability to forecast aggregate order flow. Our analysis does not enable us to be precise about this but we believe that it is likely due to Specialists and Big traders being able to process and exploit public information more rapidly than Other traders.[12] We would attribute their second advantage, in forecasting order flows, to access to high quality flows from

[12]Love and Payne (2008) provide evidence of order flow variation around macroeconomic announcements that is consistent with some traders processing information more rapidly than others.

the customer segment of the market. In the cross-rates, Big and Triangular traders, who likely see decent chunks of the customer segment of the market, demonstrate an ability to forecast order flows, but we cannot identify a class of traders in cross-rates who demonstrate superior information processing capabilities. This could be due to higher costs of identifying and exploiting information in the crosses.

5 Conclusion

This paper has introduced and analyzed a data set from a major foreign exchange trading platform. What is unique about the data set in the foreign exchange microstructure context is that it contains banker and trader identifiers. The obvious limitations of our data are the relatively short time-series length of the data we employ (one month) and the fact that we only observe a portion of any dealer's activity (that portion executed on the Electronic Broking Services platform.). The analysis provides evidence that rules out the interpretation that the widely observed price impact of foreign exchange order flow order flow is mainly a liquidity effect as is argued by, for example Berger, Chaboud, Chernenko, Howorka, and Wright (2008). By contrast, it supports the Evans and Lyons (2002) approach which highlights the information content of order flow.

We achieve this by identifying classes of dealer who have consistent and significant differences in their information quality, as revealed in the price impacts of their trading. In liquid exchange rates, dealers who specialize their activity in a given pair have the highest quality information, while in cross-rates individuals who trade the relevant currency triangle are best informed. Moreover, we show that traders located on larger trading floors have superior information to those located on smaller floors. These results give clear insight into the manner in which, and environments in which, traders glean informational advantages.

We go on to show those dealers whose trades tend to convey more information also have an advantage in forecasting order flows. This ability to forecast flows is especially pronounced for the less liquid rates in our sample. Controlling for the effect of flows on prices, we find that trades of Specialists in liquid dollar rates still carry more information than those of other types, but that the differences between trader classes in the cross-rates becomes insignificant. We speculate from these last results that our partition of traders allows us to identity subsets of dealers with differing access to private information (perhaps via customer order flows) and different capabilities in processing public information. The relationship between the release of public information and the information content of trades for our various trader classes is our next topic of study.

References

Bacchetta, P., and E. van Wincoop, 2006, "Can information heterogeneity explain the exchange rate determination puzzle?," *American Economic Review*, 96, 552–576.

Bank for International Settlements, 2007, *Triennial Central Bank Survey of Foreign Exchange and Derivative Market Activity, Preliminary Global Results.* B.I.S., Basle.

Barker, W., 2007, "The Global Foreign Exchange Market: Growth and Transformation," *Bank of Canada Review*, Autumn, 3–12.

Benveniste, L., A. Marcus, and W. Wilhelm, 1992, "Whats special about the specialist?," *Journal of Financial Economics*, 32, 61–86.

Berger, D., A. Chaboud, S. Chernenko, E. Howorka, and J. Wright, 2008, "Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data," *Journal of International Economics*, 75, 93–109.

Biais, B., P. Hillion, and C. Spatt, 1995, "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse," *Journal of Finance*, 50, 1655–1689.

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

Bjonnes, G. H., C. L. Osler, and D. Rime, 2008, "Asymmetric Information in the Interbank Foreign Exchange Market," *SSRN eLibrary.*

Bollen, N., T. Smith, and R. Whaley, 2004, "Modeling the bid/ask spread: measuring the inventory-holding premium," *Journal of Financial Economics*, 72, 97–141.

Breedon, F., and P. Vitale, 2004, "An Empirical Study of Liquidity and Information Effects of Order Flow on Exchange Rates," *SSRN eLibrary.*

Chakravarty, S., 2001, "Stealth trading: which traders' trades move stock prices?," *Journal of Financial Economics*, 61, 289–307.

Chinn, M. D., and M. J. Moore, 2009, "Private Information and a Macro Model of Exchange Rates: Evidence from a Novel Data Set," NBER Working paper 10434.

Chordia, T., R. Roll, and A. Subrahmanyam, 2001, "Order Imbalance, Liquidity, and Market Returns," *Journal of Financial Economics.*

Dufour, A., and R. Engle, 2000, "Time and the Price Impact of a Trade," *Journal of Finance*, 55, 2467–2498.

Evans, M., and R. Lyons, 2002, "Order flow and exchange rate dynamics," *Journal of Political Economy*, 110, 170–180.

Goodhart, C., R. Love, R. Payne, and D. Rime, 2002, "Analysis of spreads in the dollar/euro and deutschemark/dollar foreign exchange markets," *Economic Policy*, 17, 535–552.

Hau, H., 2001, "Location Matters: an Examination of Trading Profits," *Journal of Finance*, 56, 1959–1983.

Hau, H., W. Killeen, and M. Moore, 2002, "How has the Euro Changed the Foreign Exchange Market?," *Economic Policy*, 17, 151–177.

Ito, T., and Y. Hashimoto, 2006, "Intraday Seasonality in Activities of the Foreign Exchange Markets: Evidence from the Electronic Broking System," *Journal of Japanese International Economics*, 20, 637–664.

Killeen, W., R. Lyons, and M. Moore, 2006, "Fixed versus Flexible: Lessons from EMS Order Flow," *Journal of International Money and Finance*, 25, 551–579.

Linnainmaa, J., 2007, "Does it Matter Who Trades? Broker Identities and the Information Content of Stock Trades," *SSRN eLibrary*.

LondonFX, 2008, `http://www.londonfx.co.uk/autobrok.html`.

Love, R., and R. Payne, 2008, "Macroeconomic news, order flows and exchange rates," *Journal of Financial and Quantitative Analysis*, 43, 467–488.

Lyons, R., 1995, "Tests of Microstructural Hypotheses in the Foreign Exchange Market," *Journal of Financial Economics*, 39, 321–351.

——— , 1997, "A Simultaneous Trade Model of the Foreign Exchange Hot Potato," *Journal of International Economics*, 42, 275–298.

Lyons, R., and M. Moore, 2009, "An Information Approach to International Currencies," mimeo, University of California, Berkeley.

Osler, C., 2008, "Foreign Exchange Microstructure: a Survey of the Empirical Literature," Forthcoming, Springer Encyclopedia of Complexity and System Science.

Payne, R., 2003, "Information Transmission in Inter-dealer Foreign Exchange Transactions," *Journal of International Economics*, 61, 307–329.

Vega, C., and P. Pasquariello, 2007, "Informed and Strategic Order Flow in the Bond Markets," *Review of Financial Studies*.

Table 1: A 10 second slice of the EBS deal record

Date	Time	Rate	Side	Price	Vol.	MBank	MTrader	TBank	TTrader
08/02/99	00:01:20	EUR/USD	S	1.0668	2	03257	10003	00407	09891
08/02/99	00:01:24	USD/JPY	S	114.85	1	02633	08831	01987	10133
08/02/99	00:01:26	EUR/USD	S	1.0668	1	03257	10003	00503	10015
08/02/99	00:01:29	USD/JPY	S	114.85	2	02633	08831	01407	11623
08/02/99	00:01:29	EUR/JPY	S	122.52	4	00443	09551	00973	07493
08/02/99	00:01:30	EUR/USD	B	1.0668	3	00503	10015	01977	07585

Notes: (a) Date: Date in GMT. (b) Time: Time in GMT. (c) Rate: Currency pair. (d) Side: B for Buy, S for Sell. (e) Price: Deal price. (f) Vol.: Volume in millions. (g) MBank: Maker bank code: an arbitrary number that is unique to a bank. (h) MTrader: Maker trade code: an arbitrary number that is unique to a trader within the Maker bank. (i) TBank: Taker bank code: an arbitrary number that is unique to a bank. (j) TTrader: Taker trade code: an arbitrary number that is unique to a trader within the Taker bank.

Table 2: Exchange rate return summary statistics

Panel A : event time returns					
	Mean	STDV	Skew	Kurt	ρ_1
EUR/USD	-2.76×10^{-6}	0.0058	0.06	90.42	-0.29
USD/JPY	-1.48×10^{-5}	0.0092	0.41	56.98	-0.27
USD/CHF	1.92×10^{-5}	0.0118	0.14	166.61	-0.17
EUR/JPY	-1.32×10^{-4}	0.0211	-0.03	250.35	-0.17
EUR/CHF	8.76×10^{-6}	0.0045	-0.05	24.58	-0.2

Panel B: calendar time returns (5 minute sampling)					
	Mean	STDV	Skew	Kurt	ρ_1
EUR/USD	-1.24×10^{-4}	0.04	-0.67	14.33	0.02
USD/JPY	-7.67×10^{-4}	0.05	3.34	118.09	-0.02
USD/CHF	3.10×10^{-4}	0.04	0.53	13.72	0.03
EUR/JPY	-9.97×10^{-4}	0.06	0.41	28	-0.03
EUR/CHF	-2.47×10^{-5}	0.01	-0.1	7.93	-0.09

Notes: the table gives summary statistics for event-time and 5 minute sampled returns. 'Mean' gives the sample mean return, 'STDV' is shorthand for 'standard deviation', 'Skew' gives the coefficient of skewness and 'Kurt' the coefficient of kurtosis. ρ_1 is the first return autocorrelation.

Table 3: Number of trades and trade volume statistics

Rate	Trades (% of total)	Trades (daily)	Buy trades (daily)	Sell trades (daily)	Cash Volume (% of total)	Cash volume (daily mean)	Buy volume (daily mean)	Sell volume (daily mean)	Duration (seconds)
EUR/USD	50.58	20247.4	10161.8	10085.5	52.7	41102.6	20751.7	20350.9	4.3
USD/JPY	32.73	13102.9	6743.4	6359.5	30.1	23461.8	12001.5	11460.4	6.6
USD/CHF	8.36	3348.2	1687.1	1661.1	8.1	6293.1	3168.9	3124.1	25.6
EUR/JPY	4.87	1949.9	986.9	963.0	4.6	3611.3	1818.7	1792.6	43.7
EUR/CHF	3.45	1379.8	691.9	687.9	4.5	3476.0	1766.7	1709.4	61.4

Notes: the table gives summary statistics for the number of trade and volume for our five exchange rates. The column labelled 'Trades (% of total)' gives the percentage of the entire sample of trades in the specified exchange rate. The next three columns give the averages number of trades per day, followed by the average number of buys and sells.. The next four columns present analogous statistics for volume (rather than the number of trades). All volume statistics are measured in '000s of base currency units. The final column gives mean duration i.e. the time between trades measured in second, excluding weekends.

Table 4: Size distribution of dealt volumes

Rate	Deal size (% frequency)					Size statistics		
	1	2	3	4	≥ 5	Mean	Median	Mode
EUR/USD	55.62	23.80	10.21	3.34	7.03	1.92	1	1
USD/JPY	59.06	23.85	8.84	2.44	5.82	1.79	1	1
USD/CHF	53.21	25.91	11.49	2.61	6.79	1.88	1	1
EUR/JPY	56.76	26.79	9.73	1.83	4.90	1.75	1	1
EUR/CHF	41.31	25.50	14.86	5.34	12.99	2.38	2	1

Notes: the first five columns of the table give the percentages of all deals in a given rate that are of the specified size (in millions of units of the base currency). The final three columns gives mean, median and modal trade size for each rate.

Table 5: The number of traders per dealing room

Statistic	Value
Mean	3.94
Median	3.00
Mode	1.00
STDV	3.82
Skew	2.52
Kurt	12.32
Min	1.00
Max	30.00

Notes: the table gives summary statistics for the number of traders per dealing room. 'STDV' is shorthand for 'standard deviation', 'Skew' gives the coefficient of skewness and 'Kurt' the coefficient of kurtosis.

Table 6: The distribution of trading activity across dealing rooms

Panel A: total trading						
Traders per room	Share in total number of trades (%)					
	All	EUR/USD	USD/JPY	USD/CHF	EUR/JPY	EUR/CHF
Lowest Quartile	2.00	2.17	1.78	2.05	1.60	1.92
Second Quartile	9.23	8.39	10.88	8.34	8.65	8.78
Third Quartile	14.14	13.42	15.83	12.87	14.37	11.27
Top quartile	74.63	76.01	71.50	76.74	75.38	78.02

Panel B: individual trader activity						
Traders per room	Mean number of trades per trader					
	All	EUR/USD	USD/JPY	USD/CHF	EUR/JPY	EUR/CHF
Lowest Quartile	15.98	11.23	7.93	5.27	3.04	3.28
Second Quartile	61.18	32.49	30.75	11.49	7.11	7.59
Third Quartile	92.53	50.63	42.51	14.87	10.18	8.44
Top quartile	459.61	259.25	165.02	65.96	37.95	38.35

Panel C: mean trade size						
Traders per room	Mean trade size					
	All	EUR/USD	USD/JPY	USD/CHF	EUR/JPY	EUR/CHF
Lowest Quartile	1.52	1.55	1.45	1.47	1.37	1.86
Second Quartile	1.68	1.65	1.73	1.58	1.61	1.99
Third Quartile	1.73	1.71	1.75	1.65	1.67	2.07
Top quartile	1.94	1.99	1.82	1.96	1.79	2.48

Notes: Panel A of the table gives, for dealing rooms in the specified quartile of the room size distribution, the share of all trades in which traders from those rooms participate. Panel B gives the mean (across traders) number of trades done over the entire sample by a trader located in a dealing room from the specified quartile of the dealing room size distribution. Panel C gives the mean trade size for executions emanating from dealing rooms within the specified room size quartile.

Table 7: The distribution of traders by trade frequency

Frequency	All	EUR/CHF	EUR/JPY	USD/CHF	EUR/USD	USD/JPY
$\leq$ 5 mins	139	0	0	3	57	33
5-15 mins	438	8	12	37	184	126
15-60 mins	947	56	74	93	578	379
Hourly–daily	1097	499	757	820	1385	1385
$\geq$ daily	246	2304	2024	1914	663	944

Notes: the table gives the number of traders who trade, on average, within the frequency bands given in the row headings. The first column gives the numbers computed on the basis of trades in all rates and the following five columns give the frequenies for the five sample rates separately.

Table 8: Specialization of traders by rate

	EUR/CHF	EUR/JPY	USD/CHF	EUR/USD	USD/JPY
$\geq$ 90%	2	3	5	156	93
80–90%	0	0	5	28	29
70–80%	0	2	8	25	15
60–70%	1	3	7	19	23
50–60%	2	2	7	31	24
40–50%	4	3	10	33	36
30–40%	4	9	26	51	24
20–30%	14	21	20	49	27
10–20%	44	42	28	44	41
$\leq$ 10%	506	492	461	141	265

Notes: the table gives the number of traders who transact the percentage of their total activity specified in the row header, in the exchange rate given in the column header. Thus, for example, 9 traders execute between 30 and 40% of their total activity in EUR/JPY.

Table 9: Trader classification by rate

Trader Class	EUR/USD	USD/JPY	USD/CHF	EUR/JPY	EUR/CHF
Big	272	356	433	434	435
Specialist	134	79	4	1	2
Triangular	60	31	29	31	29
Other	2158	2041	1268	1266	787
Total	2624	2507	1734	1732	1253

Notes: the table gives the number of big, specialist (defined as big traders doing more than 90% of their trades in that rate) and triangular traders in the 5 rates. It also gives the number of traders not classified as big, specialist or triangular in each rate, and these are denoted 'Other'.

Table 10: Mean trade size by trader classification

Trader Class	EUR/USD	USD/JPY	USD/CHF	EUR/JPY	EUR/CHF
BM	1.91	1.83	2.05	1.85	2.59
SM	2.14	2.01	2.12	-	-
TM	1.65	1.50	1.92	1.57	2.21
OM	1.75	1.67	1.64	1.73	2.17
χ^2	5270.52	3069.97	1546.94	254.06	324.93
BT	1.88	1.84	1.96	1.87	2.49
ST	2.07	1.94	2.13	-	-
TT	1.74	1.57	1.93	1.66	2.40
OT	1.74	1.60	1.68	1.68	2.17
χ^2	3605.79	3228.82	820.77	237.41	174.12

Notes: the table gives mean trade size in each exchange rate of the eight trader classes generated by combining the four trader classes (Big, Specialist, Triangular and Other) with the maker versus taker distinction. The rows labelled χ^2, give test-statistics relevant to the null hypothesis that mean trade sizes for all categories of makers are identical and also that the mean trade sizes for all categories of takers are identical. These test-statistics have a $\chi^2(3)$ distribution under the null in the EUR/USD, USD/JPY and USD/CHF cases and are $\chi^2(2)$ in the EUR/JPY and EUR/CHF cases.

Table 11: The price impact of order flow: tick-time analysis (10 trade horizon)

Name	EUR/USD		USD/JPY		USD/CHF		EUR/JPY		EUR/CHF	
	Coeff	t-stat	Coeff	t-stat	Coeff	t-stat	Coeff	t-stat	Coeff	t-stat
BMBT	0.28	19.03	0.54	26.60	0.75	20.80	1.30	12.99	0.17	8.82
BMST	0.43	45.90	0.68	37.31	0.99	10.20	-	-	-	-
BMTT	0.20	8.19	0.29	6.27	0.72	13.04	1.56	13.72	0.24	9.42
BMOT	0.18	15.67	0.30	15.85	0.53	10.85	1.34	11.00	0.15	6.11
SMBT	0.27	27.16	0.55	27.70	0.71	8.13	-	-	-	-
SMST	0.38	60.72	0.69	42.13	0.90	3.67	-	-	-	-
SMTT	0.18	10.83	0.35	7.94	0.62	5.32	-	-	-	-
SMOT	0.14	18.06	0.32	16.95	0.50	4.83	-	-	-	-
TMBT	0.36	17.14	0.71	18.35	0.75	12.88	1.34	11.02	0.13	4.97
TMST	0.42	33.17	0.89	27.38	0.68	5.09	-	-	-	-
TMTT	0.20	5.47	0.75	7.39	0.67	6.77	1.70	11.61	0.16	4.95
TMOT	0.22	13.41	0.46	13.31	0.69	9.57	0.87	6.83	0.01	0.48
OMBT	0.47	44.99	0.80	45.31	1.12	27.34	1.69	19.36	0.32	14.52
OMST	0.59	85.24	0.97	64.33	1.23	12.78	-	-	-	-
OMTT	0.41	23.41	0.73	20.10	1.06	17.85	2.04	19.88	0.37	14.63
OMOT	0.37	45.27	0.60	38.26	0.99	20.04	1.59	16.52	0.24	9.99
BigRoomMaker	-0.02	-9.23	-0.04	-6.45	-0.08	-4.33	-0.12	-2.43	-0.03	-2.68
BigRoomTaker	0.02	6.51	0.04	5.11	0.09	4.60	0.13	2.88	0.02	2.49
Duration	6.29	15.72	8.61	11.08	3.61	4.02	4.92	2.85	1.45	3.93
Volatility	0.82	10.06	0.61	6.80	1.57	4.26	1.19	1.77	4.21	4.60
R^2	0.014		0.010		0.007		0.004		0.012	

Notes: the table gives coefficient estimates from regressions of tick time transaction price impacts (computed using a 10 trade horizon) on dummy variables that isolate the type of passive and active traders involved and on 4 other variables. These are the square root of the number of traders in the room on the maker side of the trade (bigRoomMaker) and the square root of the number of traders in the taker's room (bigRoomTaker), mean time between trades over the previous hour (Duration) and realized volatility. Each of these 4 variables has been demeaned. The four possible trader types are big (B), specialist (S), triangular (T) and other (O) such that the dummy variable labelled 'BMST', for example, picks out trades where the maker was a big trader and the taker a specialist. We also report t-statistics computed using Newey-West standard errors.

Table 12: Mean tick-time price impacts by trader type

Name	EUR/USD		USD/JPY		USD/CHF			EUR/JPY		EUR/CHF	
BM	0.27	32.94	0.46	30.18	0.75	22.06	BM	1.40	19.17	0.19	13.09
SM	0.24	41.98	0.48	33.55	0.68	8.61	SM	-	-	-	-
TM	0.30	24.83	0.70	23.28	0.70	13.93	TM	1.31	16.08	0.10	5.75
OM	0.46	76.16	0.77	64.01	1.10	31.44	OM	1.77	27.70	0.31	20.80
BT	0.34	42.22	0.65	45.57	0.83	26.56	BT	1.45	22.17	0.20	14.88
ST	0.45	84.75	0.81	64.89	0.95	11.37	ST	-	-	-	-
TT	0.25	18.16	0.53	16.19	0.76	16.55	TT	1.77	23.00	0.26	14.65
OT	0.23	35.71	0.42	32.88	0.68	17.79	OT	1.27	17.44	0.14	8.48
Difference tests	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
SM-OM	-0.21	27.15	-0.30	16.92	-0.42	4.82	TM-OM	-0.47	4.66	-0.21	9.00
ST-OT	0.23	27.45	0.39	22.30	0.27	2.94	TT-OT	0.50	5.10	0.12	4.79
SM-BM	-0.03	3.20	0.02	1.03	-0.06	0.76	TM-BM	-0.10	0.92	-0.09	4.00
ST-BT	0.11	11.23	0.16	8.58	0.12	1.34	TT-BT	0.32	3.46	0.05	2.37

Notes: the table gives the mean tick-time price impacts for each of the eight possible combinations of the four trader types and the maker/taker distinction. The numbers presented are simple averages across the dummy variables coefficients from Table11. Thus, for example, the EUR/USD entry labelled SM is the mean of the coefficients labelled SMBT, SMST, SMTT and SMOT in column 1 of Table 11. Other entries are defined similarly. Also presented are t statistics relevant to the null hypothesis that this mean is exactly zero. The final panel of the table displays the difference in the mean impacts for Specialist/Triangular-traders and Big and Other traders respectively, plus t-statistics for the null that these differences are exactly zero.

Table 13: Mean calendar-time price impacts by trader type

Name	EUR/USD		USD/JPY		USD/CHF			EUR/JPY		EUR/CHF	
Trader Type	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
BM	0.38	25.32	0.51	17.66	0.63	22.98	BM	1.10	27.86	0.18	21.55
SM	0.35	32.22	0.56	24.08	0.50	6.42	SM	-	-	-	-
TM	0.43	19.73	0.83	17.38	0.62	14.80	TM	1.06	20.82	0.15	15.63
OM	0.69	60.60	0.99	46.20	1.02	34.00	OM	1.49	35.89	0.29	33.35
BT	0.52	34.12	0.75	32.09	0.77	29.73	BT	1.18	29.48	0.20	24.39
ST	0.64	62.15	1.03	45.69	0.80	9.75	ST	-	-	-	-
TT	0.33	13.30	0.55	10.17	0.63	17.89	TT	1.46	29.00	0.24	23.18
OT	0.36	30.26	0.55	24.67	0.57	18.21	OT	1.01	24.01	0.17	18.06
Difference tests	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
SM-OM	-0.33	22.72	-0.43	14.29	-0.52	6.18	TM-OM	-0.43	6.73	-0.14	10.88
ST-OT	0.28	18.08	0.48	14.99	0.24	2.69	TT-OT	0.45	7.40	0.07	4.99
SM-BM	-0.03	1.83	0.05	1.38	-0.14	1.66	TM-BM	-0.04	0.65	-0.03	2.59
ST-BT	0.12	6.51	0.28	8.86	0.03	0.37	TT-BT	0.27	4.43	0.04	2.76

Notes: the table gives the mean calendar-time price impacts for each of the eight possible combinations of the four trader types and the maker/taker distinction. The numbers presented are simple averages across the dummy variables coefficients from the estimates of equation(2) Also presented are t statistics relevant to the null hypothesis that this mean is exactly zero. The final panel of the table displays the difference in the mean impacts for Specialist/Triangular-traders and Big and Other traders respectively, plus t-statistics for the null that these differences are exactly zero.

Table 14: Robustness: mean trader type impacts across impact measurement horizons

Panel A: $h = 5$

Name	EUR/USD		USD/JPY		USD/CHF			EUR/JPY		EUR/CHF	
Trader Type	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
BM	0.25	37.48	0.41	35.62	0.71	27.91	BM	1.40	24.25	0.19	16.61
SM	0.22	48.17	0.45	40.11	0.66	10.18	SM	-	-	-	-
TM	0.28	28.67	0.64	27.00	0.71	17.97	TM	1.28	19.54	0.11	7.84
OM	0.41	86.19	0.70	73.27	1.08	39.34	OM	1.70	33.90	0.31	25.66
BT	0.31	48.81	0.59	54.09	0.84	33.87	BT	1.39	26.56	0.20	17.90
ST	0.40	94.46	0.72	74.29	0.93	13.90	ST	-	-	-	-
TT	0.24	21.14	0.47	18.44	0.75	20.64	TT	1.75	28.50	0.26	18.72
OT	0.22	44.84	0.42	41.97	0.64	20.85	OT	1.24	22.34	0.15	11.25
Difference tests	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
SM-OM	-0.19	29.04	-0.25	17.49	-0.42	6.04	TM-OM	-0.43	5.35	-0.20	10.48
ST-OT	0.18	27.00	0.30	22.16	0.29	3.85	TT-OT	0.51	6.62	0.12	5.75
SM-BM	-0.03	3.44	0.04	2.36	-0.06	0.82	TM-BM	-0.12	1.48	-0.08	4.47
ST-BT	0.09	11.30	0.12	8.63	0.09	1.28	TT-BT	0.36	4.98	0.06	3.59

Panel B: $h = 50$

Name	EUR/USD		USD/JPY		USD/CHF			EUR/JPY		EUR/CHF	
Trader Type	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
BM	0.28	16.50	0.44	13.75	0.73	10.90	BM	1.54	10.57	0.12	4.83
SM	0.23	19.52	0.48	16.40	0.85	5.19	SM	-	-	-	-
TM	0.30	11.56	0.79	12.81	0.80	7.41	TM	1.31	7.66	0.15	4.53
OM	0.48	38.25	0.84	33.01	1.13	16.05	OM	1.72	12.73	0.25	9.21
BT	0.38	22.03	0.69	23.19	0.88	13.55	BT	1.50	10.80	0.20	7.85
ST	0.48	42.68	0.88	33.37	0.98	5.76	ST	-	-	-	-
TT	0.26	9.00	0.60	8.81	0.92	9.89	TT	1.77	11.32	0.17	5.32
OT	0.17	12.58	0.38	13.65	0.73	9.20	OT	1.30	8.74	0.15	4.95
Difference tests	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
SM-OM	-0.25	15.46	-0.36	9.84	-0.28	1.59	TM-OM	-0.41	1.96	-0.10	2.18
ST-OT	0.31	17.70	0.51	13.48	0.25	1.35	TT-OT	0.47	2.32	0.02	0.52
SM-BM	-0.05	2.51	0.04	1.04	0.12	0.72	TM-BM	-0.23	1.08	0.03	0.66
ST-BT	0.10	4.75	0.19	4.92	0.10	0.55	TT-BT	0.27	1.42	-0.03	0.61

Notes: the table gives the mean tick-time price impacts, for $h = 5$ in Panel A and $h = 50$ in Panel B, for each of the eight possible combinations of the four trader types and the maker/taker distinction. Impacts are derived form estimation of equation (2). Also presented are t statistics relevant to the null hypothesis that this mean is exactly zero. The final panel of the table displays the difference in the mean impacts for Specialist/Triangular-traders and Big and Other traders respectively, plus t-statistics for the null that these differences are exactly zero.

Table 15: Mean tick-time flow impacts by trader type

Name	EUR/USD		USD/JPY		USD/CHF			EUR/JPY		EUR/CHF	
Trader Type	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
BM	1.26	17.74	1.19	16.67	1.00	9.92	BM	1.95	15.99	0.46	2.67
SM	0.96	18.31	1.44	21.43	0.60	2.36	SM	-	-	-	-
TM	1.35	12.29	1.90	13.37	0.38	2.48	TM	0.62	4.05	-0.43	2.30
OM	2.48	46.07	2.42	39.12	1.79	17.67	OM	2.72	23.98	1.59	9.19
BT	1.73	24.62	1.81	26.98	1.06	10.81	BT	1.40	12.07	0.39	2.29
ST	1.80	37.18	2.18	35.97	0.52	1.96	ST	-	-	-	-
TT	1.13	8.98	1.30	8.33	1.46	11.08	TT	2.93	20.31	1.50	7.97
OT	1.39	24.97	1.66	26.48	0.74	6.44	OT	0.97	7.76	-0.27	1.49
Difference tests	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
SM-OM	-1.53	23.66	-0.98	12.20	-1.18	4.38	TM-OM	-2.10	11.99	-2.02	8.21
ST-OT	0.41	6.04	0.51	6.44	-0.22	0.77	TT-OT	1.97	11.45	1.77	6.95
SM-BM	-0.30	3.70	0.26	2.85	-0.39	1.46	TM-BM	-1.33	7.18	-0.89	3.65
ST-BT	0.07	0.82	0.37	4.46	-0.54	1.95	TT-BT	1.53	9.38	1.12	4.76

Notes: the table gives the mean tick-time impacts on cumulative order flow for each of the eight possible combinations of the four trader types and the maker/taker distinction. Means are derived form estimations of equation (3). Also presented are t statistics relevant to the null hypothesis that this mean is exactly zero. The final panel of the table displays the difference in the mean impacts for Specialist/Triangular-traders and Big and Other traders respectively, plus t-statistics for the null that these differences are exactly zero.

Table 16: Mean tick-time price impacts adjusted for flow by trader type

Name	EUR/USD		USD/JPY		USD/CHF			EUR/JPY		EUR/CHF	
Trader Type	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
BM	0.32	42.46	0.55	40.05	0.75	24.53	BM	1.36	20.67	0.17	13.19
SM	0.29	55.69	0.53	41.80	0.70	10.07	SM	-	-	-	-
TM	0.35	31.81	0.73	26.94	0.85	18.74	TM	1.76	23.85	0.16	10.21
OM	0.44	81.92	0.75	70.56	1.05	32.41	OM	1.64	29.05	0.28	21.47
BT	0.36	48.96	0.67	52.20	0.85	29.88	BT	1.62	27.32	0.21	17.39
ST	0.42	88.98	0.73	65.28	1.04	14.21		-	-	-	-
TT	0.33	27.10	0.66	22.49	0.70	16.83	TT	1.55	23.34	0.22	14.51
OT	0.28	47.51	0.51	44.34	0.76	22.23	OT	1.60	23.74	0.18	12.15
Difference tests	Mean	t-stat	Mean	t-stat	Mean	t-stat		Mean	t-stat	Mean	t-stat
SM-OM	-0.14	20.35	-0.22	14.19	-0.35	4.52	TM-OM	0.13	1.40	-0.12	5.53
ST-OT	0.15	19.68	0.22	13.72	0.27	3.36	TT-OT	-0.05	0.59	0.04	1.94
SM-BM	-0.03	2.94	-0.02	1.04	-0.05	0.72	TM-BM	0.40	4.33	-0.01	0.45
ST-BT	0.06	7.15	0.06	3.70	0.19	2.45	TT-BT	-0.07	0.89	0.01	0.43

Notes: the table gives the mean tick-time impacts on prices for each of the eight possible combinations of the four trader types and the maker/taker distinction. The regression specification controls for order-flow on the right-hand side. See equation (4) for the full specification. Also presented are t-statistics relevant to the null hypothesis that mean coefficients by trader type are exactly zero. The final panel of the table displays the difference in the mean impacts for Specialist/Triangular-traders and Big and Other traders respectively, plus t-statistics for the null that these differences are exactly zero.

Figure 1: The time-series behaviour of exchange rates: August 2001

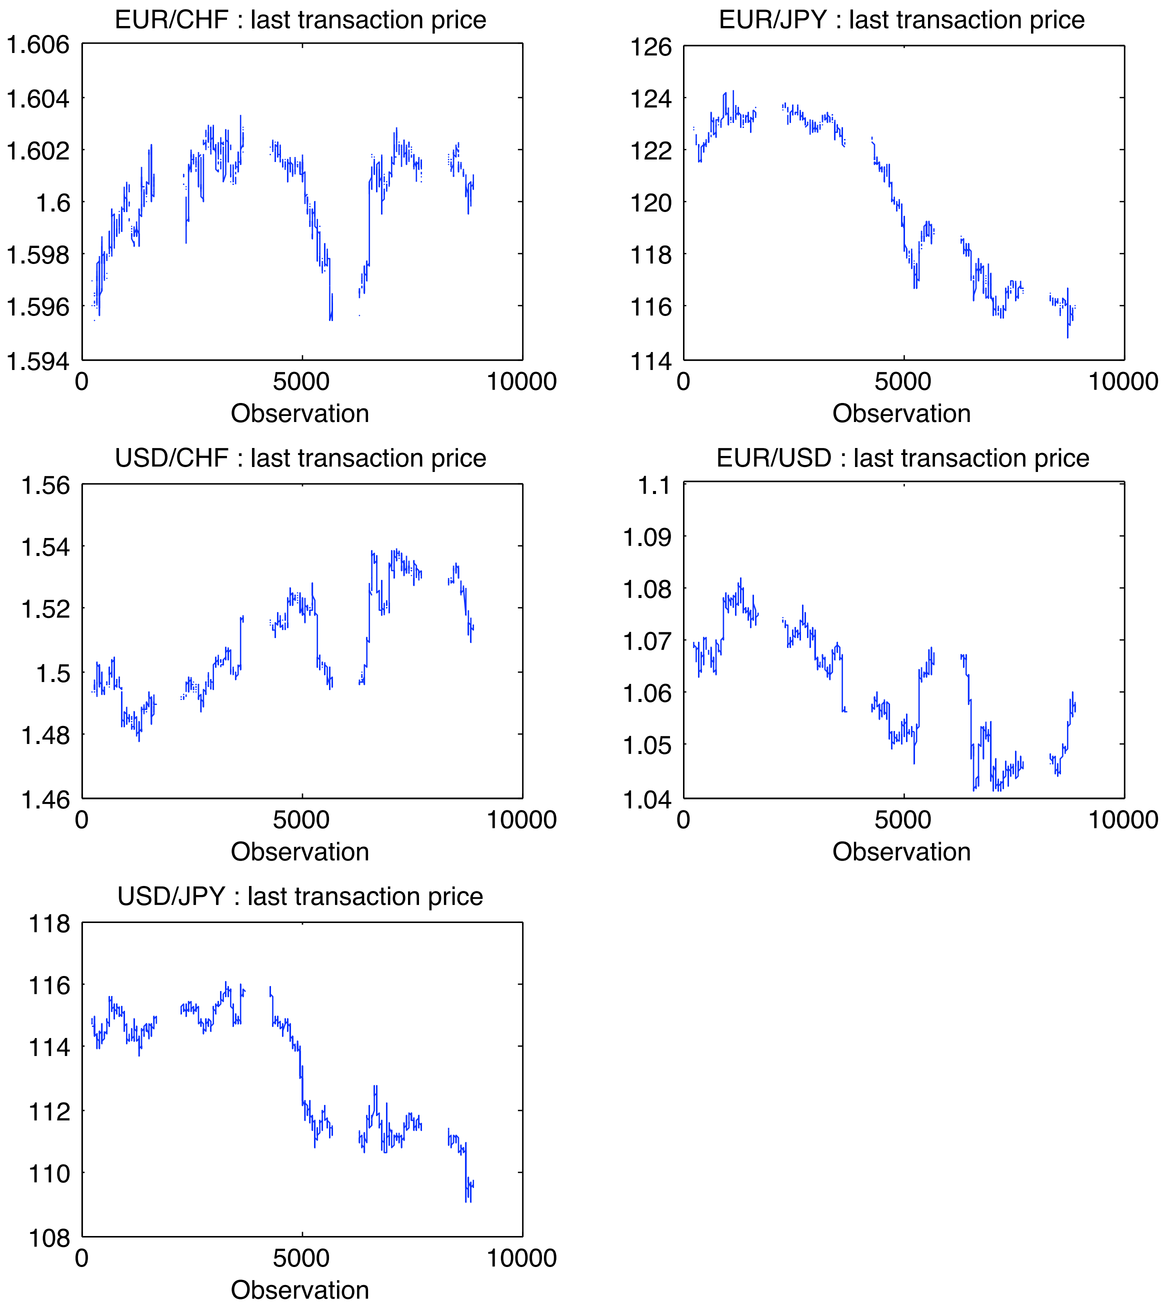

Notes: The figure shows the time-series behaviour of each of our 5 exchange rates at a 5 minute sampling frequency for the entirety of August 2001. Given the sampling frequency, each plot contains 8928 observations. When no trade in a given rate was observed in any 5 minute interval, no point is plotted. Thus the large gaps in each plotted series are generated by lack of trading on weekends.

Figure 2: Intra-day patterns in trade frequency

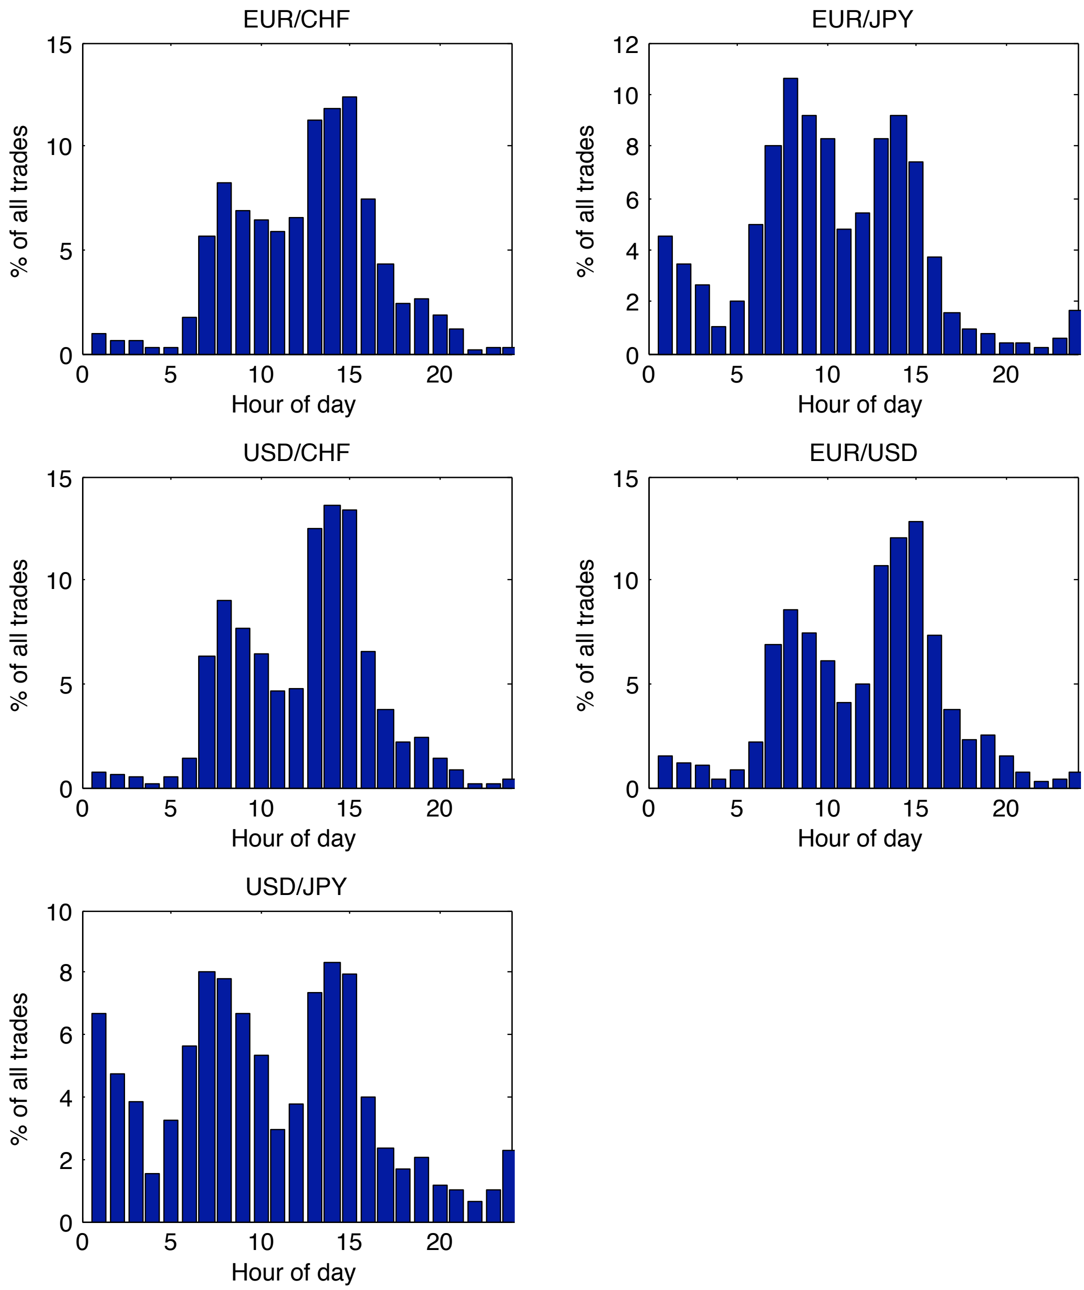

Notes: The figure shows, for each of the five sample exchange rates, the percentage of all trades executed in each hour of the trading day.

Figure 3: The distribution of the number of traders per dealing room

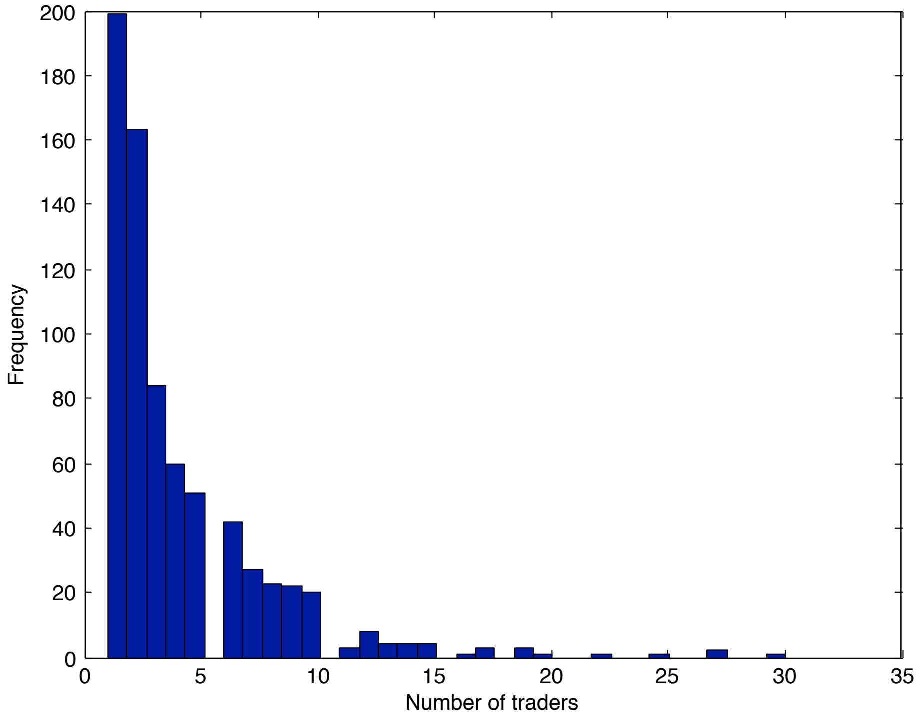

Notes: The figure shows the distribution of the number of individual traders in each trading room. Frequencies are on the y-axis and the number of traders per room is on the x-axis.