Size, specialism and the nature of informational advantage in inter-dealer foreign exchange trading

Michael J. Moore and Richard Payne*

February 14, 2009

Abstract

We present analysis of a unique foreign exchange data set which consists of a complete deal record for 1 month from the dominant foreign exchange electronic broking system complete with anonymized banker and trader identities. This is used to put flesh on the proposition that there is private information in foreign exchange markets. We find that, in liquid dollar exchange rates, traders who specialize their activity in that rate have the largest price impact from aggressive trading. In non-dollar cross pairs, the traders with the largest impact on prices are triangular arbitrage traders who spread their trades evenly across the three exchange rates in the relevant triangle. High volume traders of any kind also have reasonably large price impact from market orders. The same taxonomy of traders predicts that specialists and high volume traders suffer least price impact when supplying liquidity and this is verified in our analysis. Finally, we document the effect of a trader being located in a large institution on the information content of trades. Individuals located on large trading floors move prices further when they trade aggressively and suffer smaller adverse price movements when they trade passively.

Keywords: Foreign Exchange; Microstructure; Asymmetric Information; Order Flow

JEL classification: F31; G14

*Moore, Queens University Belfast and Harvard Center for Population and Development Studies; Payne, Warwick Business School. Correspondence: m.moore@qub.ac.uk. 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).

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 $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.

It is perhaps surprising that microstructure analysis came late to the foreign exchange markets. There are few references to the FX market in O'Hara (1995). Among the first papers to apply microstructure analysis to foreign exchange markets were Lyons (1995) and Lyons (1997). The former contains a model of FX dealer behaviour which is subsequently applied to five days of data from a single trader. The results demonstrate strong roles for both asymmetric information and inventory control effects in dealer pricing. The analysis contained in the latter paper suggests that the enormous trading volume in FX markets relative to fundamental trade and capital flows can be attributed to optimal inventory management by dealers in conditions characterized by extreme decentralization and lack of transparency. He called this 'hot potato' trading. However, the breakthrough paper in FX microstructure 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.

Economists have relatively little difficulty in accepting that order flow can impact spot FX returns through short-term liquidity or inventory effects. However, such an influence would be minor from a macroeconomic viewpoint. The real prize was to demonstrate that information itself is revealed through trading. Though it is obvious what sort of information (about future cash flows, for example) might be revealed in equity trading, it was not clear what was meant by private information in the FX context. It seemed unlikely that 'inside' information about future monetary policy or macroeconomic announcements could explain the robust results on the price impact of order flows which have subsequently proliferated in the literature. See, for example, Payne (2003).[1] In a series of papers, Evans and Lyons, both

[1] Dagfinn Rime maintains a bibliography of research related to information and order flow effects

separately and together, have developed a theoretical approach which emphasises the revelation of dispersed information about discount rates as well as cash flows. Foreign exchange trading reveals private information about idiosyncratic export demands for example: this information will eventually be aggregated and published in official statistics. However, FX order flow is not just anticipating future common knowledge information. The 'portfolio shifts' approach argues that private information about changes in risk aversion and liquidity preference, for example, will never become common knowledge and can only be impounded in price through the trading process itself. All of this suggests a rich role for the transmission of information through order flow.

Nevertheless, many economists remain reluctant to accept that order flow can have a long term impact on exchange rates. Breedon and Vitale (2004), Bacchetta and van Wincoop (2006) are pessimistic about this. Berger, Chaboud, Chernenko, Howorka, and Wright (2008) aggregate order flow for EUR/USD and USD/JPY to monthly frequencies and argue 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. Chinn and Moore (2008) respond to this robustly and show 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 an information interpretation of foreign exchange order flow, using a unique data set. The data set consists of the complete deal record for a month from a brokered inter dealer platform complete with anonymized bank and trader identifiers. 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. In addition, the larger the counter-party, the lower the price impact following the execution of its limit orders. Though

at `http://www.norges-bank.no/upload/import/english/research/dri/fxmicro-biblio.pdf`.

the analysis focuses on accounting measures of counter-party size, they provide evidence that that counter-party size is positively related to trading volume. They conclude that those institutions with the largest FX customer base have the best information.

We show that the FX market contains specialists in major liquid dollar pairs such as EUR/USD, USD/JPY, USD/CHF. These have superior information to other traders. By contrast, in non-dollar cross-pairs, such as EUR/JPY and EUR/CHF, information is concentrated among traders that specialize in arbitrage in the relevant currency triangle. We call these traders 'triangular traders'. A yen triangular trader, for example, focuses on EUR/USD, USD/JPY as well as EUR/JPY. It should be noted that the information advantage of the triangular trader is revealed in the cross-pair, not in the liquid legs of the triangle (where the triangular trader has no information advantage at all). We also find that high volume traders whether they specialize or not, show an information advantage in all markets whether dollar or cross pairs. In all markets, low volume traders display a marked information disadvantage.

Finally, we demonstrate that, for all five rates, there is a clear link between the size of of the dealing room in which a trader is situated and informational advantage. Traders located in larger rooms generate larger price impacts when trading aggressively and suffer lower adverse price movements when trading passively. These results suggest that concentration of trading professionals (and thus possibly greater research capacity and end-customer order-flow) allows the extraction of valuable information. Conversely, traders who are relatively isolated tend to have limited, if any, consistent informational advantage.

In sum, therefore, our analysis allows one to link heterogeneity in information quality across FX dealers to observable characteristics of those dealers (e.g. their location in a major bank or their specialism in a given exchange rate). As such, our results shed light on the nature and source of 'private information' in FX markets.

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 generally reliable census of foreign exchange activity in the major countries of the world.[2] 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 as EBS, for example allows hedge funds and parts of the non-dealer FX professional trading community access to its inter dealer platform. In effect, many financial institutions have now 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 foreign exchange inter dealer electronic order driven systems in the world. 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.[3] Since our data come exclusively from EBS, we concentrate on a number of its characteristic features at this point.

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

[3]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).

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. 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). There is no other information offered: nothing about size, trade direction 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 trade direction by comparing deal prices to the inside spread.

3 Data

3.1 Data structure

The data used in this study consists of the record of all deals transacted 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 brokered inter dealing trading while it was the major source (in excess of 80%) of brokered inter dealer trading in EUR/USD.[4] 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.[5] Each data record 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 raw data record is comma delimited. The first item is the date which shows as 08/02/99 – 2nd August 1999 – in each case. The second item is the time stamp which is expressed as hour: minute: second and is measured in Greenwich Mean Time. So the first deal 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 : this is reflected in the ordering of deals in the data provided to us. 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 illustrated deals, 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. Analogously, the symbol USD/JPY tells us that the rate is expressed as a Yen price for the specified quantity of Dollars.

The next part of the record gives information on the direction of trade. 'B' for a buyer initiated trade and 'S' for 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 major currency pairs. However EBS sees no commercial need to do this.[6] The next field gives the trading volume in

[4]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.

[5]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.

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

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 remaining two fields 'Taker Bank' and 'Taker Trader' relate to the aggressor, who hits the bid or lifts the offer to initiate a trade. The Bank identifiers refer to a particular trading floor so that the same bank would have different identifiers depending on whether the trade emanated from London, New York, Tokyo or another location.

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 first-order 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 negative autocorrelation disappears in calendar time (as one might expect at this relatively low frequency).

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 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 mass of trades are concentrated 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 statistics accord well with results 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 banker 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, all of Goldman Sachs' London FX activity would be grouped under one bank identifier, as would all of Deutsche Bank's Tokyo activity. The trader identifiers then isolate individuals or desks within a bank.

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 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 essentially 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. For example, 'big' trader Z might do 85% of all of his trades in EUR/USD and 15% in USD/JPY. Recalling from Table 3 that around 50% of all trades are in EUR/USD, a trader would need to be executing a very high proportion of trades in that pair to considered a 'specialist'.[7] Table 8 demonstrates a clear mass of traders in EUR/USD and USD/JPY who execute more than 90% of their trades in the specified rate 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 higher 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. In fact, one of the puzzles of the inter dealer FX market is how all three legs of a currency triangle can be simultaneously traded, if triangular arbitrage is to hold. Most of the literature on vehicle currencies, including Krugman (1980) and Rey (2001) concentrates on explaining the existence of vehicle currencies using the approximately 5 per cent of trading attributable to balance of payments flows. Lyons and Moore (2008) introduce an information approach that focuses on the other 95 percent. In essence, their model proposes that exchange rates reveal different information depending on whether trades are direct or though vehicle currencies. Arbitrage traders take account of the price impact of their trades as they exploit deviations from triangular arbitrage.

The preceding discussion leads to the concept of a 'triangular' trader, who is active in all three legs of a currency triangle. Since, from Table 4, we know that average trade sizes in each currency pair are comparable, we might caricature a yen triangular trader as one who conducts approximately 1/3 of her trades in each of EUR/USD,

[7]We think that specialist is an accurate term to use to describe these types of traders. We are aware that this term is commonly used to describe designated NYSE market makers. Such individuals provide both dealership and brokerage services and, in return for these services, the exchange grants the specialist an exclusive right to make a market for 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.

USD/JPY and EUR/JPY. An analogous remark would apply to a swiss triangular trader: approximately 1/3 of her trades should be in each of EUR/USD, USD/CHF and EUR/CHF. More flexibly, we define a triangular trader as one who conducts at least 20% of her trades in each of the three legs of a currency triangle. In terms of Table 8, this means that we exclude traders in the bottom two and top four deciles altogether. The JPY triangular traders are found at the intersection of the four deciles between 20% and 60% 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 18 and 15 respectively.

Thus four categories of trader naturally arise out of our descriptive statistics. 'Specialist' large traders in the three liquid dollar pairs; 'Triangular' large traders who are arbitrage traders in pairs including the cross; ' Big' traders are large traders that are neither Specialist nor Triangular and finally small traders that we label as 'Other'. 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 subdivide trades 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 the maker/taker distinction (the eight being BM, SM, TM, OM, BT, ST, TT and OT where it is understood that the first letter of such an abbreviation represents the trader category and the second indicates whether the trader is a maker or taker).

4 Analysis and Results

We hypothesize that the trading strategies, information quality and thus market impact of EBS participants are a function of the scale and concentration of their own order flow. It seems sensible, therefore, to focus in on the taxonomy of trader type that we have presented in Section 3.3 in order to explore information differentials.

This leads to the following hypotheses to be tested;

- In liquid pairs we expect aggressive specialist traders to exert the largest effect on prices, and specialist traders should suffer the smallest adverse price movements when trading passively.
- In the cross-rates we expect aggressive (passive) trades by triangular traders to have the largest (smallest) price impacts.
- In all rates, we expect other traders to exert the smallest price impacts when trading aggressively and to suffer the largest adverse price impacts when trading passively.

Also, we expect traders located on crowded trading floors to have access to better quality information and perhaps better quality customer order flow. Thus such traders should move prices further when they trade aggressively and suffer low price impact when supplying liquidity.

4.1 Trade Sizes

It is well known from equity markets that informed traders manage trade size. For example, Chakravarty (2001) analyses how equity traders concentrate their activity in medium sized trades as a mechanism for concealing 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.

These are presented in Table 10. 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. We interpret these results as suggesting that active management of trade size in this market in order to conceal one's trading motive is not widespread. All classes of trader use relatively similar 'normal' trade sizes 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

The most direct way to assess how relatively informed different traders are is to examine the price impact of order flow. To this end we define a trade specific price impact variable as follows:

$$\Delta p_{c,i,h,h'} = 100 \times d_{c,i} \times \ln\left(\frac{p_{c,i+h}}{p_{c,i-h'}}\right) \tag{1}$$

The subscript c indexes the currency pair; the subscript i denotes deal i and h and h' are positive integer parameters; $d_{c,i}$ is an indicator variable which is +1 for a buyer initiated trade and -1 for a seller initiated trade. So our price impact variable is the change in price, measured in basis points, from h' trades prior to h trades after deal i in currency pair c. The trade type indicator ensures that buy and sell trades are treated symmetrically. The reasons why the measure in equation (1) is used are similar to those used by Linnainmaa (2007). 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 sufficiently large, this effect is filtered out. In what follows, we set h' =5 in all cases: in other words, price change is measured from 5 trades prior to the deal of interest. Setting $h > 0$ allows the full impact of the trade on price to be felt. In the regressions below, we set h at three different values: h=5, 10 and 50. The following regression is then estimated:

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

where k and j can be Big, Specialist, Triangular and Other. σ_i is the standard deviation of the exchange rate returns, $\sqrt{DUR_i}$ is the root of the duration between trades (measured over the previous hour), $\epsilon_{c,i,h,h'}$ is an i.i.d. error term and $\beta_{i,j}, \gamma$, δ, λ_M and λ_T are parameters to be estimated. $X_{i,j}$ is an indicator variable isolating trades initiated (taken) by trader type j with liquidity supplied (made) by trader type i. For the liquid Dollar rates, there are sixteen possible trader type combinations and coefficients to be estimated and we expect all of the $\beta_{i,j}$ parameters to be positive because they are price impact coefficients from order flow. For the non-dollar cross-pairs, we have no specialist traders and thus equation (2) contains only 9 trader-type generated indicator variables (i.e. i and j can both take the values Big (B), Triangular (T) and Other (O)).

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) does this using the spread itself which is increasing in the probability of an informed trade. 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 variable which is directly proportional to $\sigma\sqrt{t}$, where σ is the standard deviation of the spot return and t is the duration between trades. In equation (2), we allow volatility and duration to enter separately. This is because Dufour and Engle (2000) argue that the price impact of trades is decreasing in the duration between trades. This goes in the opposite direction of 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 same currency pair, (measured in seconds) recorded in the hour prior to the trade.[8]

Finally, we also control for the size of the trading room in which the trader is located. Variable $MROOM_i$ is the root of the size of the room in which the maker of trade i is located and $TROOM_i$ is the size of the taker's room location.[9] 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 hypothesize that $\lambda_T > 0$ and $\lambda_M < 0$.

4.2.2 Results

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

[8]Both volatility and duration regressors have been demeaned prior to inclusion in equation (2).

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

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, passive traders in large rooms tend to suffer smaller price impacts than do passive traders in small rooms. Thus, we obtain clear evidence that the immediate environment of a trader is important in determining his (or her) effect on prices. Larger trading rooms are likely to be associated with concentration of market knowledge and research skills, and possibly bigger underlying customer bases, and 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 in basis points per trade. Across all rates, all but one of these coefficients are positive and all but two statistically significant (most very strongly so). At first sight, it may seem impracticable 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. 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). These numbers then give us the 'average 'extent to which prices drift away from the liquidity supplier after an execution. As such a large value is a bad thing from the maker's perspective.

It is clear that, in EUR/USD, both Specialist and Big Makers suffer the smallest price impacts (in that order) with the Triangular and Other Makers suffering more substantial damage following market orders. A broadly similar pattern prevails for USD/JPY. In the case of USD/CHF, Triangular liquidity suppliers do relatively well in avoiding informed takers but other than that the pattern prevails in all liquid pairs. Staying with the three liquid pairs, let us next 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 coefficient is good news for the aggressor as it indicates a rising (falling) 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 market orders in USD/JPY. In USD/CHF the Specialist and Big traders again have most impact with smaller price effects from Triangular and Other traders (though in reverse order for last two in this pair).

Thus, for the liquid rates, the orderings across making and taking are entirely consistent. When supplying liquidity, Specialists suffer smaller losses than do other trader classes, 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 a trade made by an Other trader. A similar statistic is presented for Specialist and Other takers. For all three rates, the differences between Specialists and Others are of the expected sign and highly significant.

It is similarly easy to identify the winners and losers in trading in the non-dollar cross pairs. The Triangular trader now takes on the role of the Specialist. In both EUR/JPY and EUR/CHF, Big and Triangular traders supply liquidity in a manner which avoids the price impact that Other liquidity suppliers suffer. By contrast, market orders from both Triangular and Big traders have greater impact than those from Other traders. Overall the pattern of price impact coefficients strongly indicates that the larger the order flow seen by the trader and the more specialized she is, the greater her price impact when she submits a market order and smaller the damage done to her when she supplies liquidity via limit order. Again, the test statistics in

Table 12 for equality of impacts of Triangular and Other traders, both when making and taking, add statistical weight to this claim.

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 deal time. In Tables 13 and 14, we report the results of regressions estimated in calendar time. Here the price impact variable is measured from 10 seconds before the trade (h'=10) to 60 seconds (h=60) after. 60 seconds is a long interval for a very liquid pair like EUR/USD but is about right for the less liquid cross pairs. However the pattern in coefficients is remarkably similar. For the liquid dollar pairs, the Specialist always has the biggest impact when taking liquidity, followed by the Big traders with the Triangular and Other traders far behind and close together. For the cross-pairs, the coefficient ranking when taking is Triangular (in their role as specialists), Big and Other as before. Looking at the liquidity supply coefficients in Table 12, for EUR/USD, the pattern follows the familiar ranking of largest loss for Other liquidity suppliers followed by Triangular, Big and Specialist liquidity suppliers. For USD/JPY, the Big trader performs slightly better than the Specialist with the Triangular and Other liquidity suppliers far behind in that order. The USD/CHF results are somewhat anomalous in that the Triangular liquidity suppliers avoid losses as well as the Big traders but the Specialist and Other liquidity suppliers maintain the expected extreme positions. Among the cross pairs, Other liquidity suppliers suffer noticeably larger adverse price movements than the Triangular and Big traders. The differences in mean impacts across specialist (triangular) traders and other traders for liquid (cross) rates are again all of the expected sign and extremely statistically significant.

As a final robustness check, Table 15 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 for the less liquid exchange rates. This is perhaps unsurprising, though, as on average a 50 trade horizon covers between 20, for USD/CHF and 50, for EUR/CHF, minutes for the less liquid rates (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.

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.

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, 2008, "Private Information and a Macro Model of Exchange Rates: Evidence from a Novel Data Set," *SSRN eLibrary.*

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.

Krugman, P., 1980, "Vehicle currencies and the structure of international exchange," *Journal of Money, Credit and Banking*, 12, 503–526.

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`.

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, 2008, "An Information Approach to International Currencies," mimeo, University of California, Berkeley.

O'Hara, M., 1995, *Market Microstructure Theory.* Blackwell Publishers, Oxford, U.K.

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.

Rey, H., 2001, "International trade and currency exchange," *Review of Economic Studies*, 68, 443–464.

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	EUR/JPY	S	114.85	2	02633	08831	01407	11623
08/02/99	00:01:29	USD/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	299	369	447	447	449
Specialist	134	79	4	1	2
Triangular	33	18	15	18	15
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	0.279014	18.92	0.53976	26.62	0.791894	25.41	1.30915	13.06	0.169095	11.14
BMST	0.429433	45.95	0.684195	37.33	0.897681	10.34	-	-	-	-
BMTT	0.204043	8.17	0.295141	6.29	0.585269	8.6	1.56249	13.75	0.185252	6.2
BMOT	0.178547	15.52	0.303646	15.85	0.558795	12.67	1.33705	10.96	0.122589	5.64
SMBT	0.27125	27.22	0.547644	27.72	0.694444	8.81	-	-	-	-
SMST	0.379723	60.85	0.689251	42.14	0.915745	3.73	-	-	-	-
SMTT	0.182412	10.9	0.350252	7.96	0.633978	4.07	-	-	-	-
SMOT	0.139673	18.13	0.315828	16.94	0.501171	4.83	-	-	-	-
TMBT	0.360058	17.12	0.713641	18.38	0.734771	9.37	1.35764	11.12	0.193648	5.87
TMST	0.422602	33.24	0.890651	27.41	1.02453	5.08	-	-	-	-
TMTT	0.203435	5.48	0.74915	7.4	-0.199063	-0.88	1.72189	11.71	0.233336	3.56
TMOT	0.216763	13.37	0.461688	13.33	0.65333	6.2	0.873256	6.86	0.00526503	0.14
OMBT	0.465274	44.88	0.802828	45.31	1.1354	29.62	1.69334	19.41	0.323913	16.8
OMST	0.585952	85.28	0.972704	64.33	1.24223	12.86	-	-	-	-
OMTT	0.405394	23.37	0.727686	20.12	0.926573	12.39	2.03946	19.9	0.376627	12.2
OMOT	0.365648	45.15	0.594172	38.23	0.981799	20.05	1.57944	16.48	0.240703	9.97
bigRoomMaker	-0.0244243	-9.32	-0.0407331	-6.48	-0.0788368	-4.35	-0.11944	-2.47	-0.031334	-3.32
bogRoomTaker	0.0198513	6.49	0.0363362	5.1	0.075663	4.04	0.127705	2.87	0.0262362	2.81
Duration	4.37629	16.61	5.74579	11.28	3.05397	4.44	4.47158	3.47	1.24876	4.83
rVol	0.00888257	10.79	0.00630045	6.98	0.018121	4.67	0.0147947	2.08	0.0487812	5.18
R^2	0.013		0.01		0.007		0.004		0.011	

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	
Trader Type	Mean	t-stat	Mean	t-stat	Mean	t-stat	Mean	t-stat	Mean	t-stat
BM	0.27	32.84	0.46	30.21	0.71	21.61	1.40	19.21	0.16	11.48
SM	0.24	42.09	0.48	33.56	0.69	8.22	-	-	-	-
TM	0.30	24.84	0.70	23.31	0.55	6.57	1.32	16.15	0.14	4.98
OM	0.46	76.18	0.77	64.01	1.07	29.49	1.77	27.77	0.31	19.74
BT	0.34	42.15	0.65	45.59	0.84	26.03	1.45	22.26	0.23	15.85
ST	0.45	84.79	0.81	64.90	1.02	11.18	-	-	-	-
TT	0.25	18.16	0.53	16.21	0.49	6.53	1.77	23.00	0.27	9.43
OT	0.23	35.61	0.42	32.88	0.67	15.95	1.26	17.41	0.12	6.87
Difference tests	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$
SM/TM-OM	-0.21	724.77	-0.30	285.58	-0.39	17.89	-0.45	20.48	-0.17	26.74
ST/TT-OT	0.23	766.00	0.39	498.29	0.35	11.73	0.51	27.24	0.14	17.68

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 Wald statistics relevant to the null hypothesis that this mean is exactly zero. The final two rows of the table give the difference in the mean impacts for specialist-traders/triangular-traders and other-traders when making and taking respectively, along with Wald statistics relevant to the null that this difference is exactly zero. The differences are based on specialists for the liquid USD pairs and on triangular traders for the two cross-rates.

Table 13: The price impact of order flow: calendar-time analysis (60 second horizon)

Name	EUR/USD		USD/JPY		USD/CHF		EUR/JPY		EUR/CHF	
Name	Coeff	t-stat	Coeff	t-stat	Coeff	t-stat	Coeff	t-stat	Coeff	t-stat
BMBT	0.468626	17.21	0.68549	12.99	0.720339	12.01	1.38257	12.17	0.171728	7.67
BMST	0.645764	26.9	0.947429	18.38	0.837925	9.77	-	-	-	-
BMTT	0.321092	6.18	0.273411	2.3	0.570679	8.15	1.51262	12.85	0.20141	7.79
BMOT	0.363879	15.36	0.451463	9.01	0.543313	9.15	1.1691	10.94	0.147813	6.77
SMBT	0.464707	18.84	0.719369	14.39	0.699501	8.71	-	-	-	-
SMST	0.596728	25.87	0.919259	19.61	0.617172	2.34	-	-	-	-
SMTT	0.24415	5.58	0.427601	4.87	0.362399	2.98	-	-	-	-
SMOT	0.312872	13.97	0.524182	10.65	0.401706	4.34	-	-	-	-
TMBT	0.677028	14.46	0.940904	11.69	0.748893	9.22	1.40599	10.47	0.176674	6.66
TMST	0.727945	20.3	1.27884	18.76	0.807214	4.31	-	-	-	-
TMTT	0.402461	3.56	1.11844	3.68	0.39095	2.29	1.56063	8.78	0.251128	5.83
TMOT	0.417673	11.37	0.783739	10.97	0.652032	7.11	0.769695	5.93	0.0978019	3.56
OMBT	0.781141	35.05	1.10568	23.51	1.12391	19.81	1.72679	16.09	0.300078	14.18
OMST	0.944105	45.82	1.41687	31.4	1.26703	14.62	-	-	-	-
OMTT	0.592275	14.9	0.771028	10.36	0.857463	12.57	1.84536	16.5	0.329085	13.63
OMOT	0.651808	32.9	0.924588	20.43	0.918626	16	1.63993	15.95	0.243855	11.7
SQRT BIGRoomMaker	-0.0346954	-8.08	-0.071399	-7.24	-0.0807044	-6.39	-0.152985	-6.34	-0.0277112	-5.76
SQRT BIGRoomTaker	0.0090915	2.04	0.0353802	3.61	0.0597997	4.61	0.0613395	2.65	0.0259587	5.62
Duration	2.35427	11.72	0.6507	1.46	2.11009	8.39	2.11879	4.33	0.665756	7.14
rVol	0.0146999	12	0.00735264	4.99	0.0218344	8.04	0.0213781	5.84	0.0255559	5.47
R^2	0.007		0.005		0.01		0.01		0.016	

Notes: the table gives coefficient estimates from regressions of calendar time transaction price impacts (computed using a 60 second 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 14: 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	0.45	16.58	0.59	9.72	0.67	9.91	1.35	11.64	0.17	7.33
SM	0.40	14.95	0.65	11.86	0.52	5.13	-	-	-	-
TM	0.56	14.06	1.03	10.74	0.65	7.22	1.25	9.52	0.18	6.53
OM	0.74	29.85	1.05	19.98	1.04	16.34	1.74	15.42	0.29	13.10
BT	0.60	21.76	0.86	15.77	0.82	12.12	1.51	12.64	0.22	8.97
ST	0.73	27.10	1.14	20.60	0.88	8.07	-	-	-	-
TT	0.39	8.99	0.65	6.31	0.55	6.64	1.64	12.87	0.26	9.62
OT	0.44	18.14	0.67	13.07	0.63	9.85	1.19	10.55	0.16	7.28
Difference tests	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$
SM/TM-OM	-0.34	374.23	-0.41	155.11	-0.52	36.59	-0.49	38.62	-0.12	37.95
ST/TT-OT	0.29	288.34	0.47	200.77	0.25	6.96	0.45	36.19	0.10	22.88

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 Table13. 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 13. Other entries are defined similarly. Also presented are Wald statistics relevant to the null hypothesis that this mean is exactly zero. The final two rows of the table give the difference in the mean impacts for specialist-traders/triangular-traders and other-traders when making and taking respectively, along with Wald statistics relevant to the null that this difference is exactly zero. The differences are based on specialists for the liquid USD pairs and on triangular traders for the two cross-rates.

Table 15: 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	37.36	0.41	35.66	0.69	27.30	1.40	24.35	0.17	15.06
SM	0.22	48.30	0.45	40.13	0.64	9.54	-	-	-	-
TM	0.28	28.68	0.64	27.04	0.61	9.28	1.29	19.68	0.16	6.58
OM	0.41	86.18	0.70	73.27	1.06	37.26	1.70	34.00	0.32	24.83
BT	0.31	48.73	0.59	54.12	0.84	32.45	1.40	26.70	0.22	18.29
ST	0.40	94.55	0.72	74.31	0.98	13.59	-	-	-	-
TT	0.24	21.14	0.48	18.47	0.54	9.17	1.76	28.55	0.29	12.51
OT	0.22	44.71	0.42	41.97	0.64	18.71	1.24	22.31	0.14	9.29
Difference tests	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$
SM/TM-OM	-0.18	825.19	-0.25	305.02	-0.41	31.91	-0.41	26.54	-0.16	33.54
ST/TT-OT	0.18	745.21	0.30	492.33	0.34	18.36	0.53	46.45	0.15	30.01

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	16.46	0.44	13.75	0.72	11.04	1.54	10.57	0.13	5.31
SM	0.23	19.57	0.48	16.40	0.88	5.13	-	-	-	-
TM	0.30	11.57	0.79	12.81	0.66	3.58	1.32	7.68	0.20	3.55
OM	0.48	38.24	0.84	33.01	1.09	14.62	1.72	12.75	0.25	8.73
BT	0.38	21.99	0.69	23.19	0.93	14.00	1.50	10.81	0.18	6.97
ST	0.48	42.68	0.88	33.37	1.00	5.29	-	-	-	-
TT	0.26	9.00	0.60	8.82	0.62	3.97	1.78	11.31	0.26	4.61
OT	0.17	12.53	0.38	13.65	0.80	9.02	1.30	8.74	0.15	4.20
Difference tests	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$	Mean	$\chi^2(1)$
SM/TM-OM	-0.25	235.40	-0.36	96.69	-0.21	1.27	-0.40	3.70	-0.05	0.57
ST/TT-OT	0.31	317.05	0.51	181.84	0.20	0.91	0.48	5.57	0.11	2.75

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. The numbers presented are simple averages across the dummy variables coefficients from Table13. 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 13. Other entries are defined similarly. Also presented are Wald statistics relevant to the null hypothesis that this mean is exactly zero. The final two rows of the table give the difference in the mean impacts for specialist-traders/triangular-traders and other-traders when making and taking respectively, along with Wald statistics relevant to the null that this difference is exactly zero. The differences are based on specialists for the liquid USD pairs and on triangular traders for the two cross-rates.

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

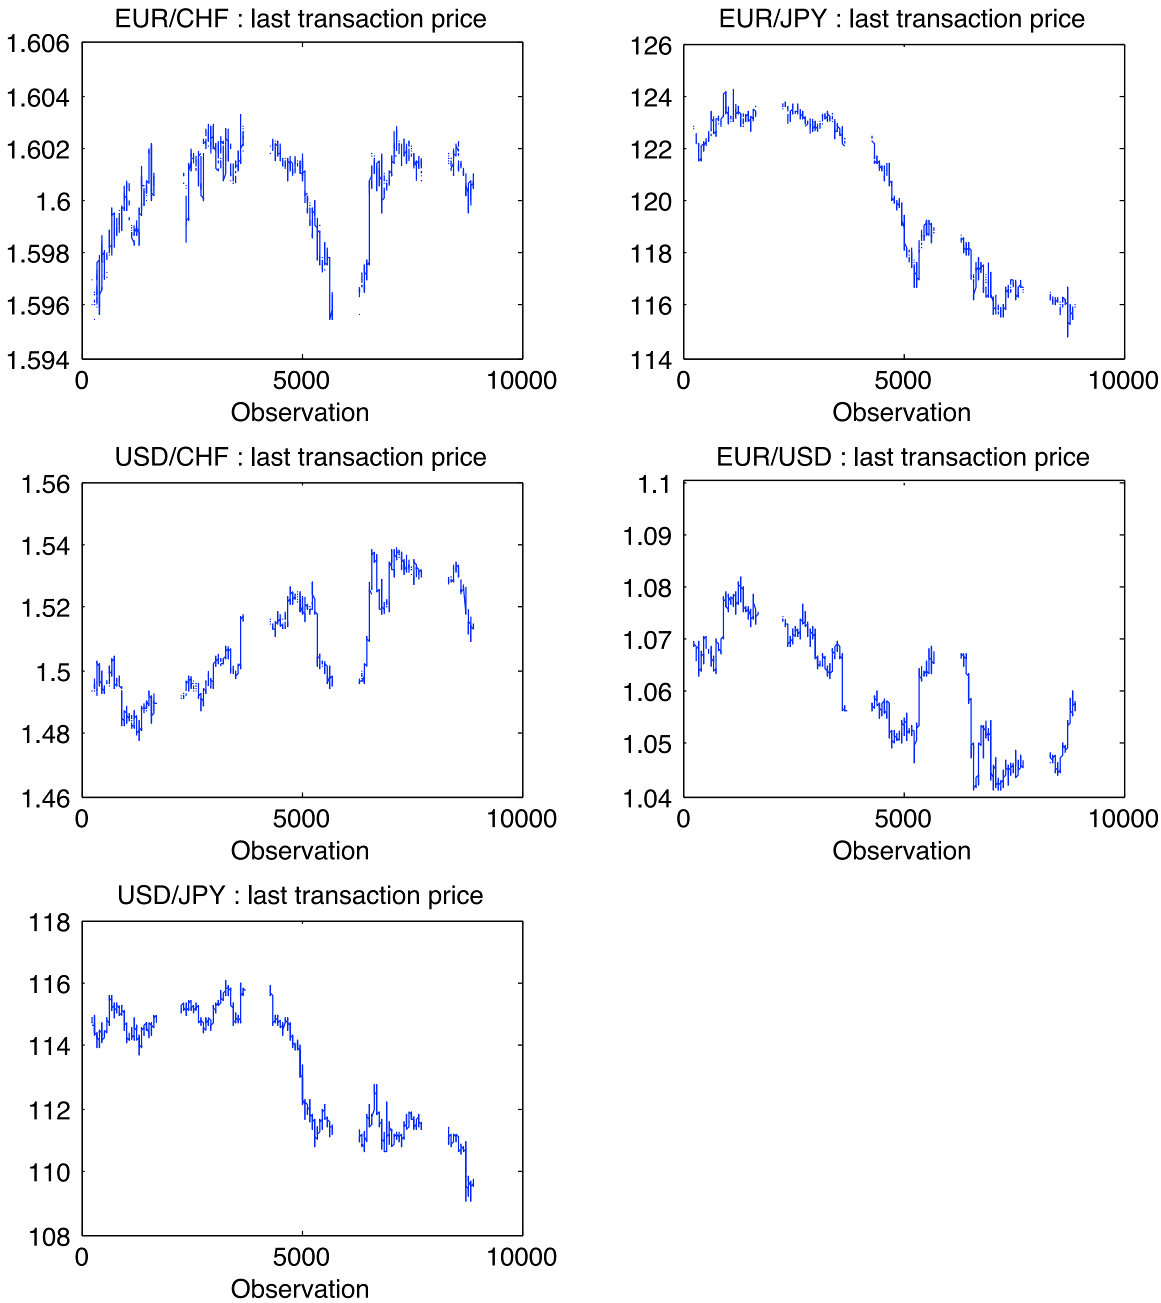

Notes: The figure shows the time-series behaviour of each of our 5 exchange rates at a 5 minute sampling frequency for the entirety of August 2001. Given the sampling frequency, each plot contains 8928 observations. When no trade in a given rate was observed in any 5 minute interval, no point is plotted.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

Figure 2: Intra-day patterns in trade frequency

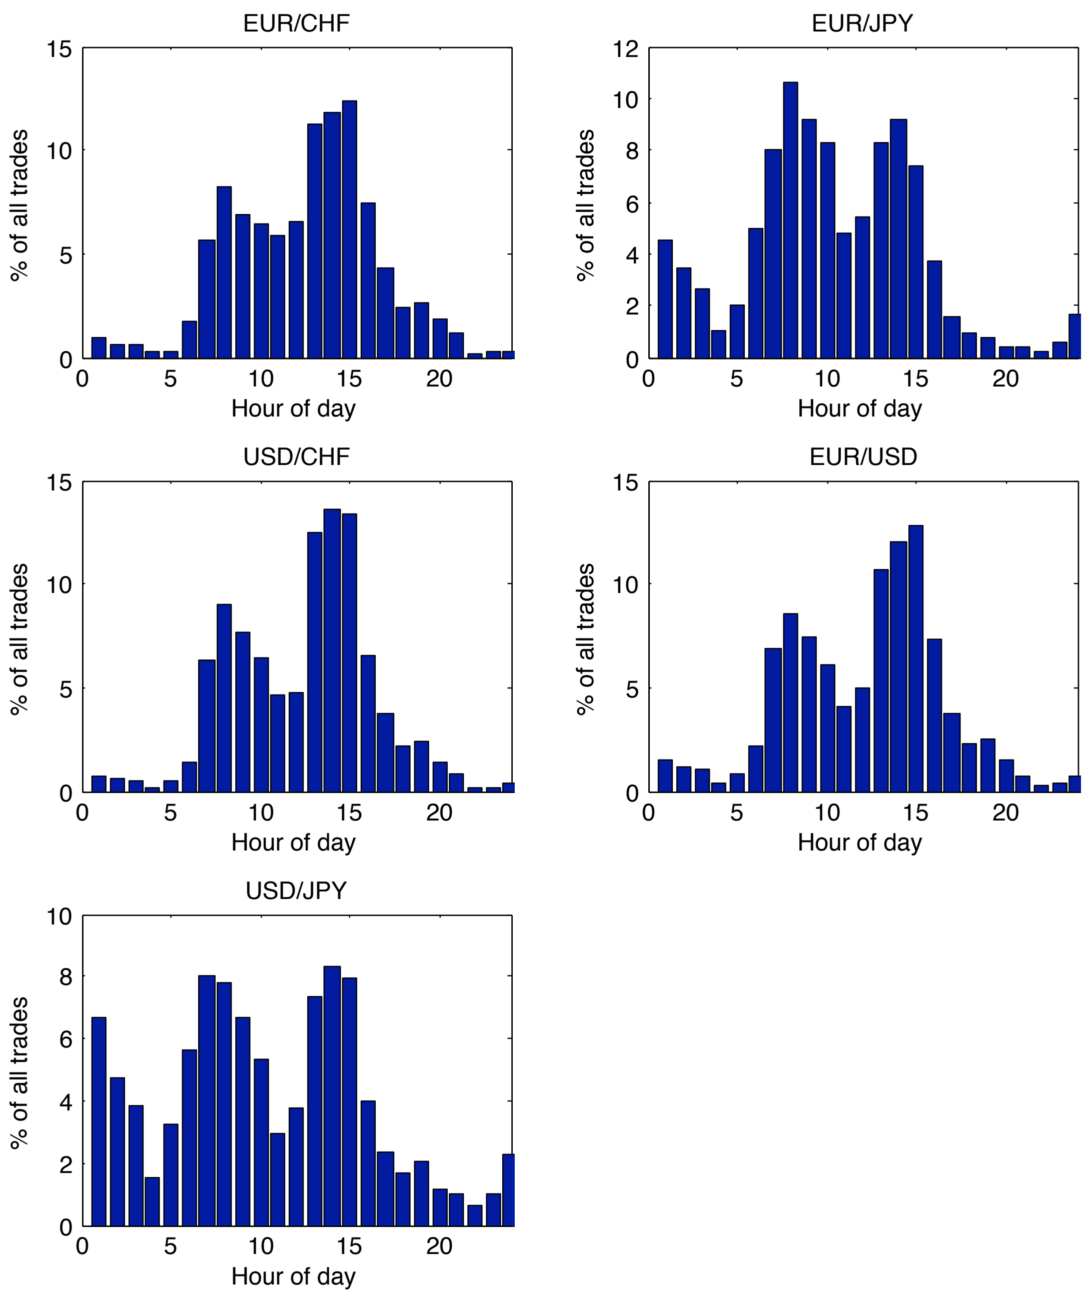

Notes: The figure shows, for each of the five sample exchange rates, the percentage of all trades executed in each hour of the trading day.

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

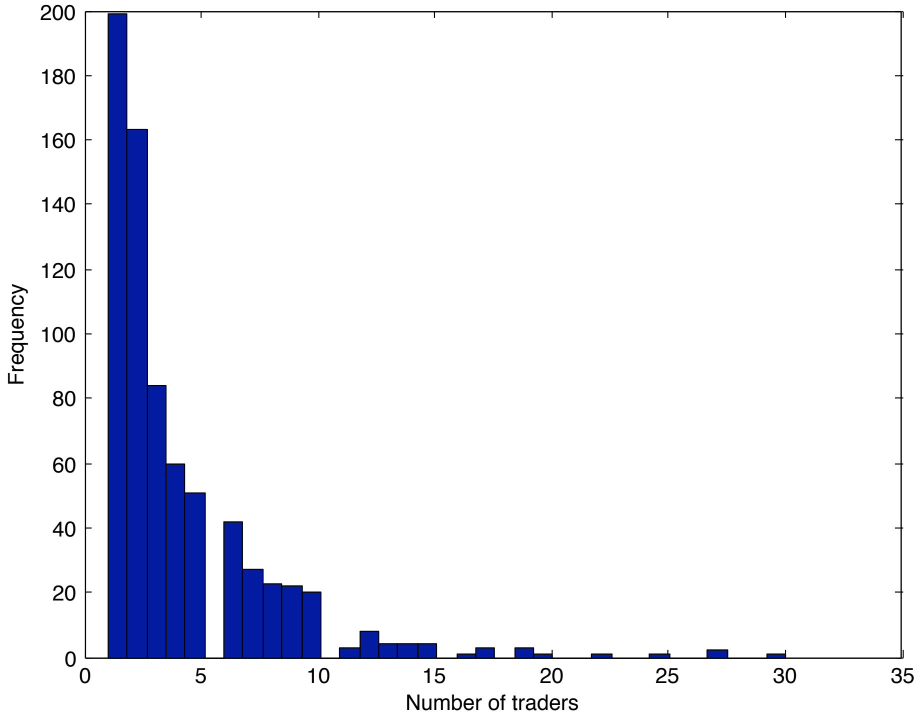

Notes: The figure shows the distribution of the number of individual traders in each trading room. Frequencies are on the y-axis and the number of traders per room is on the x-axis.