Evolutionary Reinforcement Learning in FX Order Book and Order Flow Analysis

R G Bates, M A H Dempster and Y S Romahi
Centre for Financial Research, Judge Institute of Management, University of Cambridge

Emails: rgb2@cam.ac.uk, mahd2@cus.cam.ac.uk, yr206@cam.ac.uk

Abstract: *As macroeconomic fundamentals based modelling of FX timeseries have been shown not to fit the empirical evidence at horizons of less than one year, interest has moved towards microstructure-based approaches. Order flow data has recently been receiving an increasing amount of attention in equity market analyses and thus increasingly in foreign exchange as well.*

In this paper, order flow data is coupled with order book derived indicators and we explore whether pattern recognition techniques derived from computational learning can be applied to successfully infer trading strategies on the underlying timeseries.

Due to the limited amount of data available the results are preliminary. However, the approach demonstrates promise and it is shown that using order flow and order book data is usually superior to trading on technical signals alone.

Keywords: *Reinforcement learning, evolutionary learning, FX time series analysis, order flow analysis*

1. INTRODUCTION

Published work directly relevant to the study of orders and transaction flows has appeared in relation to equity markets, but much less so for foreign exchange markets. Clearly the previous lack of data on orders and transactions in the FX market is responsible for the lack of such work, while the greater availability of data in the equity markets has produced more work of relevance (see [3]).

Early work on foreign exchange markets employed a macroeconomic approach, attempting to explain movements in exchange rates in terms of macro variables such as balance of payments, interest rate differentials, inflation differences (Purchasing Power Parity), etc. Even on timescales of a month these methods have little predictive power [16,17]. To quote Frankel & Rose [10]: "the Meese and Rogoff analysis at short horizons has never been convincingly overturned or explained. It continues to exert a pessimistic effect on the field of empirical exchange rate modeling in particular and international finance in general". Subsequent work [3,15] demonstrated that macro models begin to explain some exchange rate variation at horizons over 1 year.

In foreign exchange markets a number of studies have looked at Central Bank intervention and its effects. Dominguez [7] looks at G-3 central bank intervention, treating it as an information source. Using high frequency data the effectiveness of the intervention on the exchange rate and its volatility is assessed. The study shows that intervention is most effective when timed near major macro announcements and during periods of heavy trading volume (note not, as might be expected, during periods of low liquidity). Post-intervention mean reversion in both exchange rate and its volatility is observed.

A study of intervention by the Canadian central bank [8] attempts to understand the effectiveness of intervention by modelling the behaviour of FX dealers and uses disaggregated trades (by dealer and customer type). Net order flow (net demand or imbalance across the whole FX market) is defined as the difference between buyer-initiated and seller-initiated orders within the inter-dealer market. This definition is also used in the work of Evans & Lyons [14]. Customer initiated order flow, including Central Bank initiated flow, is found to affect volatility but its effect on the exchange rate is only short term (intra day). Whereas inter-dealer flow does have a longer term (greater than one day) effect on the exchange rate.

An information based approach to Central Bank intervention [18] looks at Bundesbank intervention. Some dealers (such as Deutsche Bank) are found to act as price leaders up to 60 minutes before intervention is publicly acknowledged. By 25 minutes prior to announcement of intervention the inter-dealer price adjustments are already two-way Granger-causal. The paper develops the view that central bank intervention (and its price effect) is revealed in stages: first to the price leader, then to (domestic) competitors and finally to the general market.

The role of information asymmetries among agents is studied in [16]. Although the FX market is global, and so we might expect spreads and volatility to be the same for all dealers irrespective of location, they find a geographical effect. At the start and end of the trading day in London and in New York there is an increase in price volatility and a widening of spreads, but only for quotes originating in the particular market opening or closing. In other words, if London and New York are examined individually they display the same pattern of spread and volatility from open to close. When New York traders start their day, with attendant high volatility and spreads, London has been trading for some hours and there is no effect on the quotes from London in either spread or volatility. Similarly the close of London has higher spreads and volatility, but there is no effect on the quotes from New York. These results are inconsistent with the standard models of asymmetric information.

An important model of the FX market involving the role of inventory and inter-dealer trading is introduced in [13,14]. The term 'hot potato trading' is introduced to describe the repeated passing of inventory imbalances between dealers. Although the imbalance is initiated by a customer order, the market state, or its willingness to accept the imbalance, is shown by the reaction in the inter-dealer market and how 'hot a potato' the imbalance proves to be. Because customer-dealer trades are not observable by the market as a whole they are not incorporated into the price until later when they are reflected in inter-dealer trading (which is observable by the rest of the inter-dealer market). Thus, in this model, it is the inter-dealer trades that drive the price.

In their work, Evans and Lyons find that order flow is important in explaining variations in nominal exchange rates over periods of about one day. Their model, which uses inter-dealer order flow accounts for about 60% of daily changes in the DEMUSD rate (in sample). They find that \$1 billion of net dollar purchases against the DM raises the DM price of a dollar by about 0.5 percent. Two definitions of order flow are used, one which measures customer-dealer flows and the other dealer-dealer flows. The customer flow figures are rapidly mean reverting (within the day) and seem to have little effect on price while the inter-dealer flows are found to mean revert very slowly. This may be a result of dealer inventory imbalances being passed between dealers over a number of days. There is some evidence (particularly in YENUSD) that buying and selling pressure is not symmetric, purchases of Yen having more effect than sales.

Perraudin & Vitale [19] also investigate inter-dealer trade and information flows. They look at the consequences for efficiency and exchange rate behaviour of the market's decentralised nature, that is, the fact that dealers are ignorant of the customer order flow of other market makers. Interbank trading is modelled as a means by which market makers "sell" each other information about their transactions with outside customers.

Finally we turn to order books: limit orders that will be executed only when the market price reaches a particular level. There seems to be no literature on order books in FX markets, probably because banks have not made this information available. What work has been done has been in equity markets, where such information has been available, particularly in Paris and London [3].

2. THE TRADING MODEL

Unlike previous literature, we attempt to approach the problem from a different perspective. Rather than attempt to explain the variation in the underlying time series, we look at whether order book and order flow data can be used to make directional bets on the movement of GBPUSD, USDEUR and USDJPY. We employ computational learning techniques as introduced in [6].

We consider agents that trade fixed position amounts in a single exchange rate. When entering a trade, the agent is able to draw on a fixed credit line from which it may borrow in either the home or the foreign currency. The money borrowed is then converted to the other currency at the current market rate thereby giving a holding of cash in one currency and a debt in the other. When the trading agent wishes to close his position he converts his cash at the new exchange rate and pays any profit into or shortfall from the account. Thus he places a series of fix-sized bets.

More formally, a trade with proportional transaction cost $\boldsymbol{c}$, exchange rates (expressed per unit of home currency) of $\boldsymbol{F_t}$ at trade entry and $\boldsymbol{F_{t'}}$ at trade exit, drawing on a credit line of $\boldsymbol{C}$ units of home currency and taking a long position in the foreign currency (and a corresponding short position in the home currency) will yield a profit of

$$C\left[\frac{F_t}{F_{t'}}(1-c)^2-1\right]. \tag{1}$$

If a short position is taken in the foreign currency (and correspondingly long in the home) then $\boldsymbol{C/F_t}$ units of foreign currency are drawn from the credit line and the profit is

$$C\left[(1-c)-\frac{F_{t'}}{F_t}\frac{1}{(1-c)}\right]. \tag{2}$$

The asymmetry of these equations is apparent and results from the profit or loss on a short position in the foreign currency being credited in the home currency. Both formulae involve transaction costs being paid per unit on two currency conversions (see [4] for a discussion of the *slippage* $\boldsymbol{c}$).

The agent considered here is able to maintain three states - long foreign (short home), long home (short foreign) or neutral (out of market).

2.1 Order Flow Indicators

Three sets of indicators are used in this work. One of the sets was derived from net daily order flow volume. This was provided to us by HSBC, the data being broken down into several different categories. The transactions were marked as to the type of client: *retail* or *institutional* and each was qualitatively judged to be *speculative* or *non-speculative*. Speculative referred to orders likely to be reversed within the day.

Filtering the data along these lines, and generating binary indicators for each category referring to whether net flow in that category was negative or positive, resulted in 9

different indicators (including unfiltered and combinations of the two categories).
The correlations between these results and the exchange rates were then calculated. This was done for the FX returns of the same day, of one day forward and of two days forward. We concentrate on these lagged results, as they would be the ones exploitable by the trading system. The trading system is fed the data with a one-day lag to ensure that the agent is trading each day on information from the day before as would be realistic in an actual daily trading framework. Due to the nature of the ERL system however, correlations in two-day lags and any other lags are also exploitable.

Table 7 illustrates the correlations for the Japanese Yen where we find the correlations of the non-lagged timeseries to be almost uniformly significant in all signals, with correlations in the region of 50%. Correlations were tested using a significance level of 90% (2-tailed test). When considering the one day lagged returns series, as expected, the number of significant correlations is reduced. The significant variables in this case were *nonspeculative trades*, *nonspeculative retail trades* and *overall retail trades*. Although the correlations in the other currencies (results not shown) were not significant at the 90% level, it is interesting that the highest correlations obtained were in these *same* categories. It appears that nonspeculative and retail trades are potentially the biggest determinants of exchange rate movements. We shall return to this point later when we consider the variables that the ERL algorithm chooses to trade with. It should be noted however that individually non-significant correlations do not preclude the ERL from inferring profitable trading strategies in terms of *combinations* of the indicators.

2.1 Order Book Indicators

An interesting aspect of the order book and one that could potentially provide the most information, was the net open orders which provides the second group of indicators.

Two types of order are distinguished: take-profit orders and stop-loss orders. Both types are only executed when the market price reaches the order price. For a purchase take-profit order, for example, the price will be below the current market price and it will be executed if the market price falls to the order price. For a purchase stop-loss order (used to protect a short position) the order will be executed when the market price rises to the order price.

For each day, the following indicators are generated:

- Net customer sales orders where the price is more than 0.0% and less than or equal to 0.5% from the current spot
- Net customer sales orders where the price is more than 0.5% and less than 1% from the current spot
- Net customer sales orders where the price is between 0 and 1% from the current spot (i.e. the sum of the former two).

These are calculated for all orders and for take-profit orders only, both for the whole order book as at the time of the snapshot of the book and for new orders only (those received in the last day). This brings the total number of indicators derived from the order book to 12.

Once again, the correlations were calculated against the FX returns series and against the return series one and two days forward. The results are shown for EURUSD in table 8. Interestingly, the correlations become significant at a lag of 2. Thus we could expect that an ERL based system being fed the indicator information lagged 1 would be able to exploit this information. We find that in all cases, the indicator generated from the net customer sales for stop-loss orders where the price is within 0.5% of the current spot gave the least amount of relevant information. The farther the stop loss prices were from the current spot, the higher the correlations. With regards to take profit orders vs. all orders, it was difficult to distinguish which could potentially yield more information. Similarly the correlations of new orders vs the ones related to the entire order book did not appear to differ in the sense of one containing more information than the other about the returns series. We return to this point again when we consider the indicators chosen by the ERL system.

The net volumes were converted into binary indicators, as before signifying whether the net volume in that category was negative or positive.

2.2 Technical Indicators

We consider a set of technical indicators to be used for our benchmark tests. These are in line with indicators used previously such as in Dempster & Jones [4] and Dempster & Romahi [6]. We employ commonly used indicators including Price Channel Breakout, Relative Strength Index, Stochastics, Moving Average Convergence/Divergence, Moving Average Crossover, Momentum Oscillator and Commodity Channel Index with parameters as suggested in [1]. In several of the indicators, we use several different parameters.

3. THE TRADING STRATEGY

We can consider the *market state* $\boldsymbol{s}$, represented by the indicator signals to be a *vector stochastic process* $\boldsymbol{F}$ driven by the *exchange rate* process and the corresponding order flow and order book and make the required trading decisions by solving the *stochastic optimization problem* defined by the maximization of expected return over the *trading horizon* $\boldsymbol{T}$ net of transactions costs, *viz.*

$$\mathbb{E}\sum_{i=1}^{\mathbf{N}_T} r_i(\mathbf{F}_{t_i}, \mathbf{F}_{t_i'}), \tag{3}$$

where $\mathbf{N}_{\boldsymbol{T}}$ denotes the random number of trades to the horizon each with return $r(\mathbf{F}_t,\mathbf{F}_{t'})$ in the home currency.

The system we consider attempts to find *approximate* solutions to this problem. It attempts to discover a *trading* *strategy* $\phi : S \times \{l,s\} \rightarrow \{l,s\}$ that maps the current market state s_t and current position (long, short or neutral) to a new position (long, short or neutral). It should be noted that although our trading strategies ϕ are formally *Markovian* (feedback rules), the technical indicators require a number of periods of previous values of rates to compute the corresponding 0-1 entries in $\mathbf{s}_t$, as does the order book.

The objective of the trading strategies developed is thus to maximize the expected home currency (dollar) return (after transaction costs) using the model described above.

3.1 Evaluation

Since we do not have an explicit probabilistic model of how exchange rates evolve, we adopt the familiar approach of dividing our data series into an *in-sample* region, over which we optimize the performance of a candidate trading strategy, and an *out-of-sample* region, where the strategy is ultimately tested.

4. EVOLUTIONARY REINFORCMENT LEARNING

The goal of reinforcement learning based trading systems is to optimize some relevant measure of trading system performance such as profit, economic utility or risk-adjusted return.

The *evolutionary reinforcement learning* (ERL) approach incorporates a reinforcement learning subsystem that solves the in-sample Markov decision problem fed by an evolutionary algorithm (EA) that attempts to avoid the overfitting tendency often exhibited by reinforcement learning. This is achieved by constraining the in-sample search space.

The role of the EA here is to choose some optimal subset of the underlying indicators that the RL system will then use. The form of EA utilised is the binary string genetic algorithm due to Holland [12]. Each bit in the bitstring represents whether or not the corresponding indicator is being fed into the RL.

With regards to fitness evaluation, the in-sample period was broken down into 2 months of true in-sample data with a further 1 month of data in the *evaluation* period which is used to evaluate individuals within the GA's population of potential solutions. The return over this second period is used as the fitness function of the GA. Once the optimal bitstring is found, the subset of indicators that the bitstring represents is fed into the RL system described next (see Figure 1).

4.1 Reinforcement Learning Subsystem

For a detailed review of the algorithm, refer to [6]. A brief overview is given below.

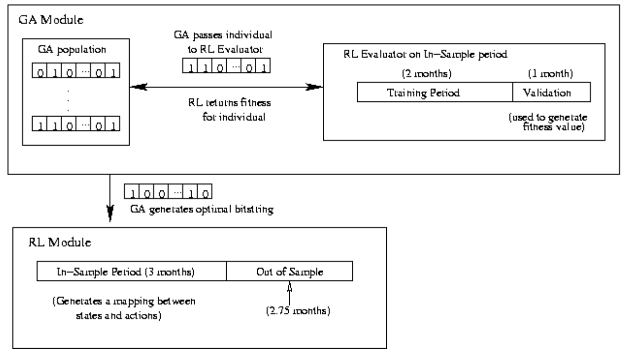

Figure 1. ERL System

Reinforcement learning systems consist of an *agent* interacting with an *environment*. At each time step t the agent *perceives* the state of the environment $\mathbf{s}_t \in \boldsymbol{S}$ and chooses an *action* $a_t \in \boldsymbol{A}$ from the set of available actions in state s_t. As a consequence of this action the agent observes the new state of the environment s_{t+1} and receives a *reward* r_t. This can be defined as a *dynamic programming* problem where the objective is to find the *policy* π (state to action mapping) that maximises the *Q-value* Q*, the value of each action a that can be taken from that state. This is given by

$$Q^*(s_t, a_t) \Leftarrow \mathbb{E}\left\{ r_{t+1} + \gamma \max_{a'} Q^*(s_{t+1}, a') \mid s_t = s, a_t = a \right\}, \quad (4)$$

where γ is the *discount factor* representing the preference given to immediate over future rewards.

We use Watkins's *Q*-learning algorithm [21] that estimates the *Q*-value function using data from the previous learning episode. The *Q*-learning *update* is the backward recursion

$$Q(s_t, a_t) \Leftarrow Q(s_{t_c}, a_{t_c}) + \left\{ r_{t+1} + \gamma \max_{a'} Q^*(s_{t+1}, a') \mid s_t = s, a_t = a \right\}, \quad (5)$$

where the current *state-action pair* $(s_t, a_t) := (s_{tc}, a_{tc})$, that from the previous learning episode. At each *iteration* (episode) of the learning algorithm, the action-value pairs associated with all the states are updated and over a large number of iterations their values converge to the optimal [20].

For our trader the state $\mathbf{s}_t$ is the market state as defined by the previously described indicators and the set of actions $\mathbf{A}$ refer to whether to take a long, short or neutral position.

5. PRELIMINARY TRADING RESULTS

In this section, we examine the performance obtained out-of-sample feeding different subsets of information to the ERL system. In order to draw conclusions from these preliminary results, we shall discuss the results of each currency in turn.

5.1 EURUSD

We first consider EURUSD. Tables 1 through 6 in the appendix demonstrate that order flow information based trading, at average out-of-sample monthly returns of 1.27% and 0.27% with slippage values of 2 and 4 bp respectively, outperforms technical indicator based trading, with returns of 0.47% and 0%. When we then feed both datasets into the ERL subsystem, we find that the results improve further: we obtain monthly average returns as high as 1.58% even at transaction costs as high as 10bp.

Similarly, we find that order book based trading is superior to the technical indicator based benchmark. The order book based trading (although it made a slight loss at 4bp), made a monthly return of 1.28% at 8bp and returned 0.45% at 10bp.

When we examined the order book indicators that were chosen by the GA for EURUSD (as shown in Table 9), we find significant consistency. More encouragingly, the highly correlated indicators (typically where the stop-losses are further from the spot as alluded to earlier) are consistently chosen.

Combining the technicals with the order book data proved to be beneficial, producing a monthly return of 1.52% at 8bp and 0.8% at 10bp (although again making a loss at 4bp but was profitable at 0 and 2bp).

The results were mixed when combining the order flow and order book data but generally encouraging (returns of 0.32%, -0.72%, 1.13% at slippage values of 0bp, 2bp and 4bp respectively) from Table 6.

We find that the ERL system makes a trade on the EURUSD on average once every 3 days, (see Tables 1 through 6).

When we consider the actual order flow indicators chosen in each case (Table 10), we find that for EURUSD, there does not appear to be consistency in the indicators used for different slippage values. However, a low correlation between the individual indicators does not imply that the ERL will be unable to exploit patterns in the data. Indeed with low correlations, we do not expect to see consistency in the choice of indicators.

Figure 1 shows the cumulative return obtained from EURUSD with 4bp slippage. It clearly shows how overlaying the technicals with order flow and book derived information improves the trading system performance.

5.2 GBPUSD

For GBPUSD (Figures 6-8), we find a broadly similar result. Although the benchmark technicals performed well in isolation (returns of 0.82%, 0.36%, 2.01% and 0.16% at slippage values of 0, 2, 4 and 8 respectively), we still find that order book data based trading was superior (though not consistently so).

When we consider the actual indicators chosen for the order book trades, we find that once again there is consistency in the choice of indicators (Table 9) – and also consistency with the EURUSD results. This is encouraging as it increases confidence in these preliminary results.

By merging the two indicator sets and feeding the combined set into the system, we are able to further improve the results and obtain returns of 1.25%, 1.68%, 0.4% and 0%)

It is important to note that once again merging the order book data with the technicals proved to be the most promising combination.

With regards to the flow data based trading however, the out-of-sample returns are mixed. Significantly, GBPUSD was also the currency where we found the correlations with the returns time series to be lowest amongst the flow correlations. When looking at the indicators chosen by the flow experiments, we do find that a consistency pattern emerges across the slippage values with *retail*, *nonspeculative institutional* and *speculative institutional* being consistently selected.

5.3 USDJPY

Finally, with regards to the USDJPY, the ERL system interestingly, chose to remain out of the market over the in-sample period examined. This is an important result in itself as it illustrates that in situations where there may be a low signal to noise ratio, rather than 'memorize' the noise (as a pure RL system would be inclined to do), the ERL system remains out of the market. When considering the indicators chosen, however, we do find that the ERL system managed to consistently choose the highly correlated indicators, namely *nonspeculative retail*. Unlike the other currencies, the correlations obtained in this case were highly significant (at the 95% level). However, it is usually the case that the ERL system often remains out of the market if only one or two indicators are chosen to be fed into the RL subsystem. This would need to be considered further on a longer data set as it appears that there could potentially be more information here than the ERL is currently exploiting.

Runtime for the training periods tended to be less than 10 minutes when looking at any set of indicators in isolation. When we merge the indicator sets, we increase the search space by a factor 2^n where n is the number of added indicators. The search space rapidly increases with n. We found that with over 20 indicators, the optimisation routine ends prematurely. In the first instance, the GA evolved for 500 generations. This was increased to 20000 generations and the initial population size was also increased to 100. This increased computation time to 7 hours.

6. CONCLUSIONS AND FURTHER WORK

Having considered the results, it is clear that the analysis of order book and order flow data for automated FX trading is a promising area for further investigation. Naturally, due to the limited time horizon of the data examined, this work can only be termed a preliminary analysis. It should be noted that the reinforcement learning algorithm is of order O(n) thus feeding it with double the amount of data only doubles the runtime. There should be no problems therefore in scaling-up the algorithm to larger data sets. Indeed, our previous work with technical indicators considered 15 minute intraday trading and trained over an entire year of such high frequency data.

Care must be taken in providing the algorithm with more degrees of freedom. The GA is of order $O(2^n)$ in inputs. Thus for every additional input variable, the searchspace doubles in size. The current input variables were fed in as binary indicators. While this was promising, we will no doubt be able to improve the performance by increasing the resolution of the indicators. This would significantly increase the searchspace for the evolutionary algorithm. At present, the training process of the algorithm takes 10 minutes to run. As was alluded to earlier, evolutionary algorithms in general belong to the class of algorithms termed *embarrassingly parallel computations*, so that close to linear speedup can be achieved through parallelisation. By choosing the input variables intelligently or incorporating heuristics into the GA, we can also constrain the state space explosion.

Further refinements that are being investigated in this work are the use of risk adjusted objectives and the overlay of cash/risk management techniques to limit losing trades.

Table 1. Average monthly returns on benchmark technical indicator tests (Number of monthly trades)

Currency	Slippage	In-sample[1]	Out-of-Sample[2]
EURUSD	0	1.46 % (19)	0.09 % (10)
EURUSD	2	1.02 % (21)	0.47 % (11)
EURUSD	4	0 % (0)	0 % (0)
EURUSD	8	0 % (0)	0 % (0)
EURUSD	10	0 % (0)	0 % (0)
GBPUSD	0	1.14 % (16)	0.82 % (15)
GBPUSD	2	1.64 % (2)	0.36 % (2)
GBPUSD	4	1.72 % (2)	2.02 % (1)
GBPUSD	8	1.16 % (2)	0.16 % (2)
GBPUSD	10	0 % (0	0 % (0)
USDJPY	0	0 % (0	0 % (0)
USDJPY	2	3.23 % (28)	1.54 % (16)
USDJPY	4	0 % (0)	0 % (0)
USDJPY	8	0.14 % (1)	0.94 % (1)
USDJPY	10	0 % (0	0 % (0)

Table 2. Average monthly returns on order book tests (Number of monthly trades)

Currency	Slippage	In-sample	Out-of-Sample
EURUSD	0	2.11 % (11)	1.86 % (9)
EURUSD	2	2.68 % (11)	0.88 % (8)
EURUSD	4	2.12 % (7)	-0.13 % (12)
EURUSD	8	2.19 % (6)	1.28 % (5)
EURUSD	10	2.89 % (5)	0.45 % (4)
GBPUSD	0	1.60 % (8)	3.02 % (8)
GBPUSD	2	1.50 % (9)	-0.29 % (13)
GBPUSD	4	0.87 % (6)	0.37 % (6)
GBPUSD	8	0.95 % (1)	1.52 % (0.4)
GBPUSD	10	0.89 % (2)	1.64 % (4)

[1] In-sample period runs from the 1st of March, 2002 until the 30th of May, 2002 (65 data points)

[2] Out-of-sample period runs from the 1st of June, 2002 until the 19th of August, 2002 (58 data points)

Table 3. Average monthly returns on order flow tests (Number of monthly trades)

Currency	Slippage	In-sample	Out-of-Sample
EURUSD	0	2.69 % (14)	-1.09 % (15)
EURUSD	2	3.06 % (9)	1.15 % (8)
EURUSD	4	2.78 % (7)	0.27 % (6)
EURUSD	8	2.04 % (10)	0.20 % (9)
EURUSD	10	2.86 % (11)	-0.23 % (8)
GBPUSD	0	2.68 % (12)	-0.13 % (14)
GBPUSD	2	2.24 % (11)	0.42 % (12)
GBPUSD	4	2.42 % (10)	0.16 % (5)
GBPUSD	8	1.91 % (8)	-0.13 % (8)
GBPUSD	10	1.59 % (6)	0.20 % (4)
USDJPY	0	0 % (0)	0 % (0)
USDJPY	2	0 % (0)	0 % (0)
USDJPY	4	0 % (0)	0 % (0)
USDJPY	8	1.07 % (2)	-0.28 % (0.3)
USDJPY	10	2.70 % (7)	-0.59 % (7)

Table 4. Average monthly returns on joint technical and order flow rests (Number of monthly trades)

Currency	Slippage	In-sample	Out-of-Sample
EURUSD	0	3.40 % (11)	1.48 % (11)
EURUSD	2	2.94 % (11)	0.18 % (13)
EURUSD	4	3.02 % (11)	1.23 % (9)
EURUSD	8	3.95 % (11)	-0.63 % (3)
EURUSD	10	3.75 % (6)	1.58 % (3)
GBPUSD	0	2.68 % (12)	-0.13 % (14)
GBPUSD	2	1.66 % (12)	-0.52 % (12)
GBPUSD	4	1.32 % (8)	1.12 % (6)
GBPUSD	8	0.83 % (7)	-0.47 % (8)
GBPUSD	10	1.15 % (3)	1.04 % (3)
USDJPY	0	2.90 % (14)	0.97 % (11)
USDJPY	2	3.44 % (10)	2.64 % (6)
USDJPY	4	2.11 % (13)	1.32 % (12)
USDJPY	8	3.58 % (8)	2.26 % (9)
USDJPY	10	2.68 % (9)	1.93 % (8)

Table 5. Average monthly returns on joint technical and order book tests (Number of monthly trades)

Currency	Slippage	In-sample	Out-of-Sample
EURUSD	0	3.22 % (10)	0.17 % (13)
EURUSD	2	2.26 % (6)	1.17 % (5)
EURUSD	4	1.57 % (11)	-0.75 % (11)
EURUSD	8	3.66 % (6)	1.52 % (6)
EURUSD	10	2.38 % (4)	0.81 % (3)
GBPUSD	0	1.05 % (5)	1.25 % (6)
GBPUSD	2	0.97 % (8)	1.68 % (10)
GBPUSD	4	1.10 % (4)	0.40 % (5)
GBPUSD	8	0 % (0)	0 % (0)
GBPUSD	10	0 % (0)	0 % (0)

Table 6. Average monthly returns on joint order flow and order book tests (Number of monthly trades)

Currency	Slippage	In-sample	Out-of-Sample
EURUSD	0	2.60 % (10)	0.32 % (9)
EURUSD	2	1.64 % (9)	-0.72 % (11)
EURUSD	4	2.38 % (5)	1.13 % (7)
EURUSD	8	2.60 % (8)	-0.15 % (10)
EURUSD	10	2.07 % (5)	0.72 % (7)
GBPUSD	0	2.36 % (11)	0.83 % (7)
GBPUSD	2	0 % (0)	0 % (0)
GBPUSD	4	0 % (0)	0 % (0)
GBPUSD	8	0.67 % (3)	1.81 % (2)
GBPUSD	10	1.05 % (4)	1.77 % (4)

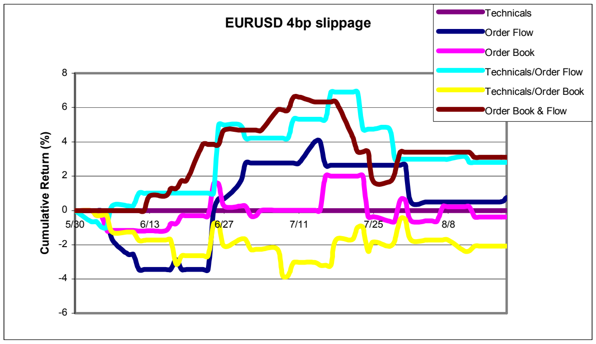

Figure 1. Out-of-sample cumulative return EURUSD 4bp

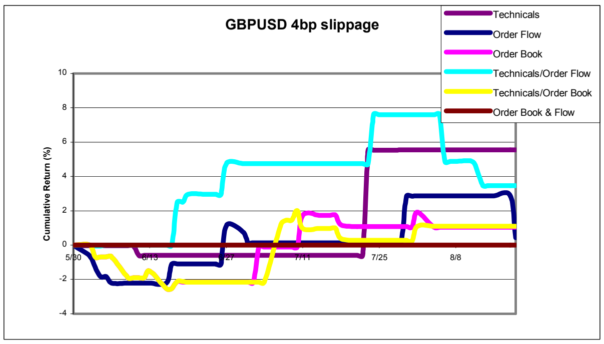

Figure 2. Out-of-sample cumulative return GBPUSD 4bp

Table 7. Order flow indicator correlations: USDJPY[3]

	FXReturns	lag(1)	Lag(2)
Unfiltered	***52.999 (6.87)***	3.902 (0.42)	9.746 (1.06)
Retail Trades	9.581 (1.05)	***30.416 (3.49)***	11.091 (1.21)
Speculative Retail Trades	-4.187 (-0.46)	-4.501 (-0.49)	-0.073 (0)
Nonspeculative Retail Trades	6.593 (0.72)	***28.793 (3.29)***	12.449 (1.36)
Institutional Trades	***54.004 (7.05)***	-0.529 (-0.05)	9.488 (1.03)
Speculative Inst. Trades	***47.187 (5.88)***	-4.204 (-0.46)	0.97 (0.1)
Nonspeculative Inst. Trades	***37.35 (4.42)***	8.445 (0.92)	**17.902 (1.98)**
Non-Speculative Trades	***35.831 (4.22)***	**16.156 (1.79)**	***19.777 (2.2)***
Speculative Trades	***46.596 (5.79)***	-4.94 (-0.54)	0.741 (0.08)

Table 8. Order book indicator correlations: EURUSD

Net Customer Sales	FXReturns	lag(1)	lag(2)
New: 0%-0.5% of spot	13.502 (1.38)	14.2 (1.44)	-5.17 (-0.52)
New: 0.5%-1% of spot	10.729 (1.09)	9.209 (0.93)	-13.178 (-1.33)
New: 0%-1% of spot	***19.407 (2)***	**18.871 (1.94)**	-14.16 (-1.43)
New: 0%-0.5% of spot (TP)	3.411 (0.34)	10.077 (1.02)	-3.62 (-0.36)
New: 0.5%-1% of spot (TP)	-4.102 (-0.41)	0.335 (0.03)	***24.358 (2.52)***
New: 0%-1% of spot (TP)	-0.988 (-0.1)	7.07 (0.71)	**17.094 (1.74)**
All: 0%-0.5% of spot	12.573 (1.28)	8.606 (0.87)	-6.34 (-0.63)
All: 0.5%-1% of spot	5.752 (0.58)	0.21 (0.02)	***-19.556 (-2)***
All: 0%-1% of spot	15.201 (1.56)	7.624 (0.77)	***-20.098 (-2.06)***
All: 0%-0.5% of spot (TP)	2.652 (0.26)	12.175 (1.23)	-1.847 (-0.18)
All: 0.5%-1% of spot (TP)	-3.19 (-0.32)	3.992 (0.4)	***27.469 (2.87)***
All: 0%-1% of spot (TP)	-0.383 (-0.03)	11.799 (1.2)	***18.633 (1.9)***

Table 9. Indicators Used in the Order Book Experiments[4]

Currency/ Slippage	New: OB1[5]	New: OB2[6]	New: OB3	New: OB2 (TP[7])	New: OB3[8] (TP)	All: OB1	All: OB2
Eur 0			x				X
Eur 2			x				X
Eur 4	x	x			x		x
Eur 8			x		x	x	
Eur 10		x	x		x	x	x
Gbp 0	x	x		x	x		x
Gbp 2	x	x		x	x		x
Gbp 4		x	x	x		x	x
Gbp 8							x
Gbp 10							x

Table 10. Indicators Used in the Order Book Experiments

Currency/ Slippage	All	Ret[9]	Ret (NS[10])	Ret (S)	Inst	Inst (NS)	Inst (S)	NS	S
Eur 0					x				x
Eur 2						x			x
Eur 4	x					x		x	
Eur 8	x	x	x		x	x	x		
Eur 10	x	x	x		x		x	x	x
Jpy 0			x						
Jpy 2			x						
Jpy 4			x						
Jpy 8		x	x	x					
Jpy 10	x	x	x	x		x	x		
Gbp 0	x	x	x		x	x	x		
Gbp 2	x	x			x	x	x	x	
Gbp 4	x	x	x	x		x	x		
Gbp 8	x	x		x		x			
Gbp 10	x	x				x	x	x	

[3] Correlations are bolded when significant at the 90% level (2 tailed test) and are bolded and italicized when significant at the 95% level (p-values shown in parentheses)

[4] Note that all indicators that were consistently not used were not included in the table

[5] OB1 refers to sales for stop loss orders where the price is within 0.5% of the current spot

[6] OB2 refers to sales for stop loss orders where the price is between 0.5%-1% of the current spot

[7] TP refers to "Take Profit" orders only

[8] OB3 refers to sales for stop loss orders where the price is between 0%-1% of the current spot

[9] Ret refers to retail and Inst to institutional trades

[10] NS refers to non-speculative while S refers to speculative

ACKNOWLEDGMENT

This work was partially supported by HSBC plc. The authors would like to thank Mark Austin, Stacey Williams and Adrian Hughes of HSBC Investment Bank for their collaboration and encouragement.

REFERENCES

[1] S Achelis, *Technical analysis from A to Z*, McGraw Hill, 2000.

[2] J-P Bouchaud, M Mezard and M Potters, Statistical properties of the stock order books: empirical results and models *Quantitative Finance* **2** 251-6, 2002.

[3] M Chinn & R Meese, Banking on Currency forecasts, *J. International Economics* **38** 161-178, 1994.

[4] MAH Dempster, CM Jones, A real-time adaptive trading system using genetic programming, *Quantitative Finance* **1** (July 2001) 397-413.

[5] MAH Dempster, TW Payne, GWP Thompson & YS Romahi, Computational Learning Techniques for FX Trading Using Popular Technical Indicators, *IEEE Transactions on Neural Networks*, Special issue on Computational Finance, March 2001, **12**, pp 744-754.

[6] MAH Dempster & YS Romahi, Intraday FX Trading: A hybrid credit assignment evolutionary reinforcement learning approach, H.Yin et al. (eds): IDEAL 2002, *Lecture Notes in Computer Science* **2412**, pp. 347-358, 2002. Springer-Verlag Berlin Heidelberg.

[7] KM Dominguez, The market microstructure of Central Bank intervention, National Bureau of Economic Research, Working Paper 7337, 1999.

[8] C D'Souza, A market microstructure analysis of foreign exchange intervention in Canada, Bank of Canada Working Paper 2002-16.

[9] MDD Evans & RK Lyons, Order flow and exchange rate dynamics, *J. Political Economy* **110** 170-80, 2002.

[10] J Frankel & A Rose, Empirical research on nominal exchange rates, In *Handbook of international economics*, Eds G Grossman & K Rogoff, Elsevier Science 1689-1729, 1995.

[11] DA Hsieh & AW Kleidon, Bid-Ask spreads in foreign exchange markets: Implications for models of asymmetric information. In *The microstructure of foreign exchange markets*, Ed JA Frankel, G Galli & A Giovammini, U. of Chicago Press, 1996.

[12] JH Holland, *Adaptation in natural and artificial systems*, University of Michigan Press, Ann Arbor, MI, 1975.

[13] RK Lyons, Foreign Exchange: Macro puzzles, micro tools, Working Paper Haas School, UC Berkeley, 2001. (Forthcoming in *Economic Review*.)

[14] RK Lyons, *The microstructure approach to exchange rates*, MIT Press, 2001.

[15] N Mark, Exchange rates and fundamentals: Evidence on long-horizon predictability, *American Economic Review* **85** 201-218, 1995.

[16] R Meese & K Rogoff, Empirical exchange rate models of the seventies, *J. International Economics* **14** 3-24, 1983.

[17] R Meese & K Rogoff, The out-of-sample failure of empirical exchange rate models. In *Exchange rate and international macroeconomics*, Ed J Frenkel, U. of Chicago Press, 1983.

[18] B Peiers, Informed traders, intervention, and price leadership: A deeper view of the microstructure of the foreign exchange market, *J. Finance* **52** 1589, 1997.

[19] W Perraudin & P Vitale, Interdealer trade and information flows in a decentralized foreign exchange market. In *The microstructure of foreign exchange markets*, Ed JA Frankel, G Galli & A Giovammini, U. of Chicago Press, 1996.

[20] R Sutton & AG Barto, *Reinforcement learning: An introduction*, MIT Press, Cambridge MA, 1998.

[21] C Watkins, Learning from Delayed Reward, PhD Thesis, King's College, University of Cambridge, 1989.

[22] I Zovko and J Doyne Farmer, The power of patience: a behavioural regularity in limit-order placement *Quantitative Finance* **2** 387-392, 2002.