

Adaptive systems for foreign exchange trading

Mark P Austin¹, Graham Bates², Michael A H Dempster^{2,3}, Vasco Leemans²
and Stacy N Williams¹

¹ HSBC Global Markets, 8 Canada Square, London E14 5HQ, UK

² Centre for Financial Research, Judge Institute of Management, University of Cambridge, Cambridge CB2 1AG, and Cambridge Systems Associates Ltd., 1 Earl Street, Cambridge CB1 1JR, UK

E-mail: mahd@cus.cam.ac.uk

Received 22 July 2004, in final form 12 August 2004

Published 3 December 2004

Online at www.tandf.co.uk/journals/titles/14697688.asp DOI: 10.1080/14697680400008593



1. Introduction

Foreign exchange markets are notoriously difficult to predict. For many years academics and practitioners alike have tried to build trading models, but history has not been kind to their efforts. Consistently predicting FX markets has seemed like an impossible goal but recent advances in financial research now suggest otherwise. With newly developed computational techniques and newly available data, the development of successful trading models is looking possible. The Centre for Financial Research (CFR) at Cambridge University's Judge Institute of Management has been researching

trading techniques in foreign exchange markets for a number of years. Over the last 18 months a joint project with HSBC Global Markets has looked at how the bank's proprietary information on customer order flow and on the customer limit order book can be used to enhance the profitability of technical trading systems in FX markets. Here we give an overview of that research and report our results.

2. Macro models don't work

It has long been known that macroeconomic models fail to predict exchange rates at time horizons less than 12 months (Meese and Rogoff 1983, 1997). Our interest has been in the techniques used by FX traders and market

³ Author to whom any correspondence should be addressed.

makers and has focussed on predicting intraday or daily exchange rates. Because economic variables (including interest rates) have little explanatory power at these frequencies many traders have used technical analysis. Technical analysis attempts to predict markets by identifying patterns in the price action and is in common use among practitioners. Indeed this has become the dominant forecasting method used by intraday traders and is the area on which we have concentrated.

Despite its century long history amongst investment practitioners, technical analysis methods have tended to be regarded with scepticism by academics. Furthermore, users of technical analysis had never made serious attempts to test the predictions of the various techniques they use. However, as evidence has accumulated that markets are less efficient than academics had believed (for example, LeBaron 1999) there has been a resurgence of academic interest in the area. For examples of academic investigations into the performance of technical analysis indicators, see Allen and Karjalainen (1999) and Neely and Weller (2003).

The work of the Centre for Financial Research in this area is summarized in various publications: Jones (1999), Dempster and Jones (2001, 2002), Dempster *et al.* (2001), Dempster and Romahi (2002) and Bates *et al.* (2003).

Much of the work has concentrated on high frequency (intraday) trading but more recent work making use of customer transaction flows and limit order book data has been into ways of optimising daily rather than intraday trading. This recent work is also reported in detail in two research dissertations: Romahi (2003) and Leemans (2003).

Some of our early work studied specific chart patterns such as the widely used ‘Channel’ or ‘Head and Shoulders’ patterns, while latterly we have investigated methods of developing optimal trading rules by combining a number of technical or informational indicators. It is this data selected combination of a range of indicators which looks most promising.

Formally, we consider the *market state* \mathbf{s} represented by the indicator signals to be a *vector stochastic process* driven by the *exchange rate* process \mathbf{F} and make the required trading decisions by solving the *stochastic optimization problem* defined by the maximization of expected return over the *trading horizon* T net of transaction costs:

$$\mathbb{E} \sum_{i=1}^{N_T} r_i(\mathbf{F}_t, \mathbf{F}_t'),$$

where N_T denotes the random number of trades to the horizon with each 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* $\phi: S \times \{l, s, n\} \rightarrow \{l, s, n\}$ that maps the current market state \mathbf{s}_t and the 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 previous values of exchange rates to compute the corresponding entries in \mathbf{s}_t . The objective of the trading strategies developed is thus to maximize the expected home currency return, after transaction costs.

3. Machine trading using technical indicators alone

We have been developing techniques that use a number of technical indicators in combination to give an optimal trading rule for a particular currency pair. Such a trading system then gives an optimal trade-off between risk and return. An important goal was also to discover ‘understandable’ trading rules. As such, care has been taken to create indicators based on rules commonly adopted in the market. A trivial example of this would be e.g., buy if the RSI is greater than 0.7 AND the Stochastics is not indicating overbought. RSI and Stochastics are technical signals routinely used by traders and are available in most technical charting packages. It must be stressed however that making sensible indicators is relatively straightforward, compared with the task of identifying genuinely useful combinations of them. The last problem is very challenging since as the number of indicators is increased the number of possible combinations tends to explode. This has limited the scope of earlier work which has often looked at indicators individually and it is in combining the signals that significant advances have been made. Much of our work has used either a Genetic Algorithm (GA) approach or the techniques of Evolutionary Reinforcement Learning (ERL), both of which are described in the papers cited above. Both the GA and ERL approaches provide an understandable selection and combination of indicators, which is in contradistinction to some other machine learning techniques such as neural networks where it is not possible to ‘unpick’ the trading rule; the result of a neural network has to be viewed as a black box. The GA and ERL techniques we have adopted are much more transparent and yet still extremely powerful in problems as complex as this. The general structure of our automated trading system is represented in figure 1.

The results, using only technical indicators, have shown consistent profitability (Table 1). Using FX market data from 1994 to 1998 at 15 min frequency on Yen–Dollar, Sterling–Dollar and Swiss Franc–Dollar, trading was profitable on all three currency pairs, out of sample, even after allowing for costs of up to 2 basis points for the opening and closing of transactions (4 basis points overall). With some machine learning techniques

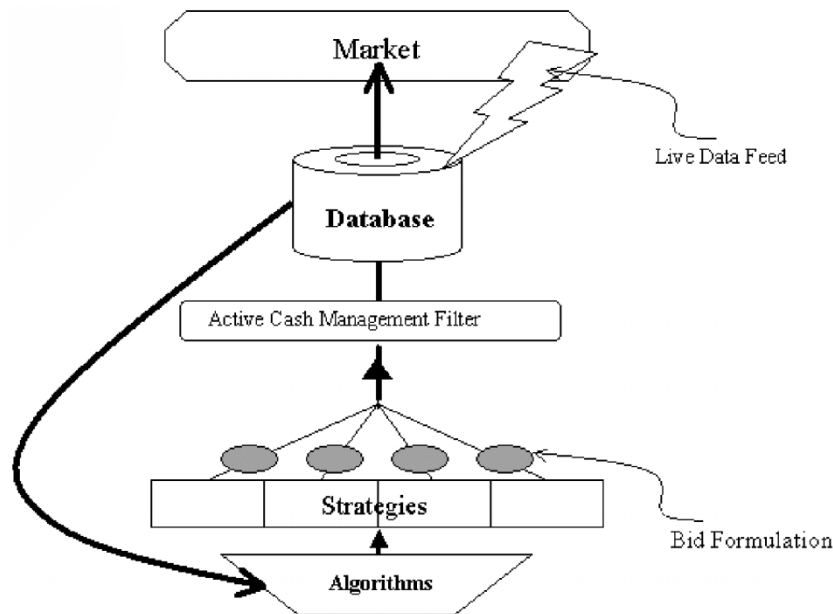


Figure 1. General structure of our trading systems.

Table 1. Out-of-sample trading returns using technical indicators alone.

Currency	Slippage (bp)	Annualized return (%)	Sharpe ratio	Significance level
GBPUSD	0	54.88	1.82	0.01%
GBPUSD	1	20.81	1.62	0.01%
GBPUSD	2	13.07	0.76	10%
GBPUSD	4	-1.81	-0.21	N/S
GBPUSD	8	-1.59	-0.07	N/S
GBPUSD	10	0	0	N/S
USDCHF	0	80.79	2.08	0.01%
USDCHF	1	12.53	0.33	5%
USDCHF	2	4.29	0.24	10%
USDCHF	4	2.42	0.18	N/S
USDCHF	8	-0.77	-0.14	N/S
USDCHF	10	1.43	0.07	N/S
USDJPY	0	91.93	1.36	0.01%
USDJPY	1	18.79	0.34	10%
USDJPY	2	-0.37	0	N/S
USDJPY	4	-6.44	-0.08	N/S
USDJPY	8	-1.36	-0.17	N/S
USDJPY	10	-4.25	-0.08	N/S

profitability extends to costs of 8 basis points overall in some currencies.

Using more recent data on Euro–Dollar from January 1999 to January 2002, at 1 min frequency, similar results were obtained using a new algorithm that was fed with technical indicators and past returns only. However the trading rules were profitable out of sample only with costs less than or equal to 2 basis points per opening or closing transaction (figure 2). These spreads are realistic for inter-bank trading as spreads on the main electronic dealing system used, the Electronic Broking System (EBS), vary inversely with trading volume over

the 24 day between one and five basis points. Results are similar for Cable (Dollar–Sterling), although the algorithm required a longer initial period to learn how to trade this market (figure 3). This is a clear sign that foreign exchange markets have become more efficient over the last few years. The Euro market in particular is the largest and most liquid market, and therefore one of the most efficient markets in its pricing. This is also a likely explanation for the flattening of the profit curve in figure 2 as time progresses. Although a combination of technical indicators may have been profitable in the past we must now ask if other

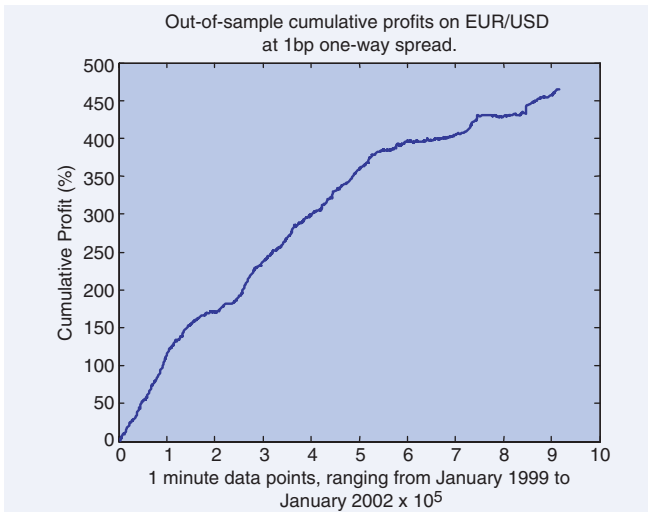


Figure 2. Out-of-sample cumulative trading returns for EUR/USD at 1 bp one way spread (equals 2 bp two way spread).

information can be used to supplement the technical indicators to improve trading performance in current FX markets.

4. Why might flow and order data be useful?

While the results using only technical indicators are encouraging, other sources of information are considered for inclusion into the trading rules. In particular it is investigated whether non-public information available only to an FX market maker could contribute to profitability of the trading strategy.

As one of the major foreign exchange market makers, HSBC receives *market orders* (for immediate execution) and *price limit orders* (which are executed when the market moves to a certain level) from a wide range of customer types: long-term investors, corporates and short-term speculators such as hedge funds. The question is if this information could be used to help predict the direction of various FX markets. The influence of this information on trading performance was analysed, both when used in isolation and in combination with technical indicators.

In our analysis we refer to market transactions which are to be executed immediately as *flows* and to price limited transactions which are executed only when the market price reaches a specified level as *orders*.

The flow data consisted of daily totals of executed transactions for the period March 2002 to February 2003 in the currencies: Dollar–Yen, Euro–Dollar, Sterling–Dollar and Euro–Sterling. The orders are a snapshot of all open orders taken at the same time each day for these same currencies.

It is evident that flow and order book data can help in forecasting markets, considering what information each can give us. Flows show what different types of clients are doing. For example: if speculative institutional traders, such as Hedge Funds, are buying Euros and selling Dollars we reach a very different conclusion about the future direction of Euro–Dollar than if the buying of Euros came from a US importer of French wine who had bills to pay in Euros and had little choice in the timing of his purchase of Euros. Transactions from certain types of customer provide more useful information about the future direction of the market than transactions from other types of customer.

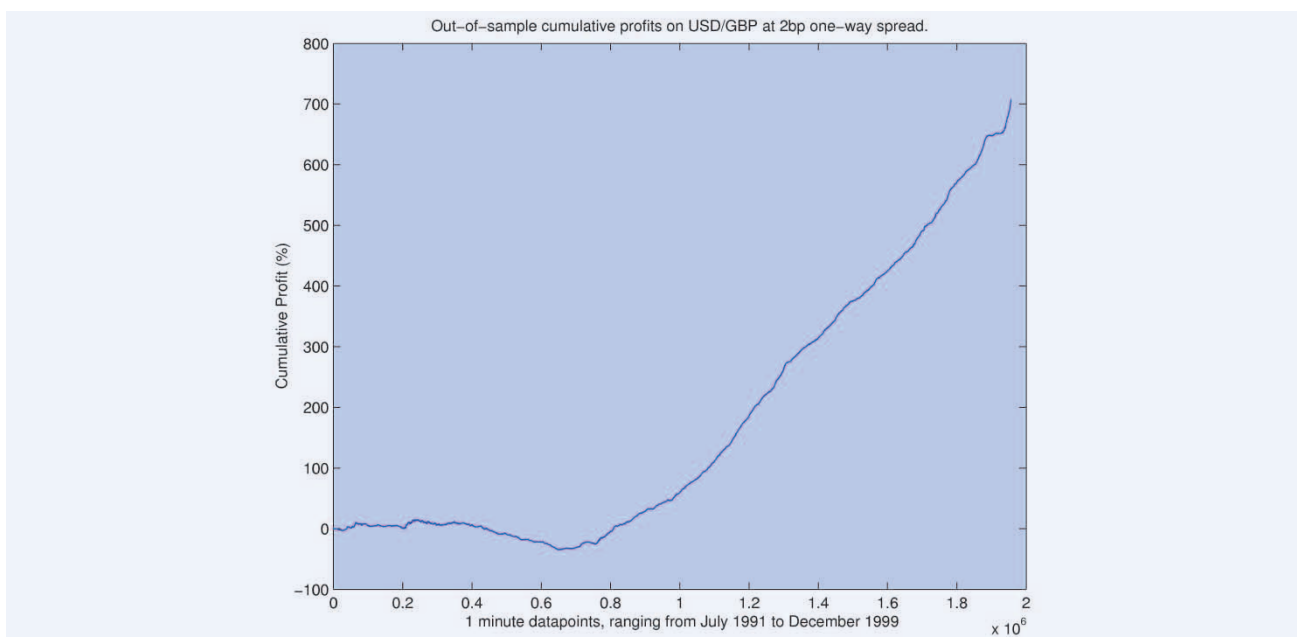


Figure 3. Out-of-sample cumulative trading returns for USD/GBP at 1 bp one way spread (equals 2 bp two way spread).

The order book also provides interesting information as it gives the open customer orders for a major market maker and, as such, is representative of the latent demand or supply pressure on the market. In particular it indicates whether moves in the exchange rate will trigger reinforcing or dampening transactions as the market orders are executed. For the initial analysis six months of order book information were used, from March 2002 to August 2002. It should be noted that the order book gives a snapshot of potential (buying or selling) pressures on the market. It consists of transactions waiting to happen when the market price reaches a particular level.

In the order book two types of orders are distinguished: *take profit* orders and *stop loss* orders. These have very different effects when they are activated by a price move in the market. A *take profit* order acts in the opposite direction to the market move that triggers it, a *stop loss* order acts in the same direction as the market move that triggers it. An example provides an illustration of this point.

Example

A customer has bought dollars at a price of 100 Yen; if the price rises to 110 there will be a profit. If the customer leaves a limit order to sell the dollars at a price of 115 Yen this would be a **take profit** order. When the price *rises* the order is to *sell*, that is the resulting transaction acts in the opposite direction to the move that triggers it.

If the customer had again bought dollars at a price of 100 Yen but the price was now 90 a limit order could be left to close the position—that is to sell the dollars—if the price fell to 85. This would stop the loss getting any bigger and is thus known as a **stop-loss** order. In this case a *fall* in the market price triggers a *sale* and so acts in the same direction as the move that triggers it.

Although these examples are for sale transactions, stop loss and take profit orders can be for purchases in the case where the original transaction was a short sale. Similarly limit orders can be used to open a position as well as to close it. The same terminology is used when the limit order opens a position: a take profit order acts in the opposite direction to the market move that triggers it. A rise in price triggers a take profit sale to open a short position and a fall in price triggers a take profit purchase to open a long position.

It is clear that take profit and stop loss orders have potentially very different effects on the market and must be carefully distinguished. One reinforces market moves, the other dampens down market moves.

5. Statistical results are encouraging

Although our ultimate objective was to blend this data into the machine learning systems described earlier, we

found it useful to undertake some linear statistical analysis first. This preliminary analysis should be indicative of how useful the proprietary flow and order data might be. Measurable linear structure would be an encouraging sign for the more complex nonlinear machine learning techniques that are applied in a second phase. Statistical analysis of the contemporaneous relationship between FX order flow and price has recently been performed by Lyons (2001) and evidence was found that flow does impact markets and by varying amounts depending on the type of customer involved.

The transaction flows from the HSBC archive were divided into four categories depending on the type of client as follows:

- speculative or leveraged investors (such as Hedge Funds),
- institutional investors (such as pension funds, usually acting through asset management firms),
- corporates (trade or long-term capital flows),
- others, including central banks.

For each day two flow numbers were calculated for each of these categories: a gross and a net flow. The gross flow is the total buy transaction volume plus total sell volume, which gives an indication of the volume of business from that type of client each day. The second was the net flow, buying volume minus selling volume, which gives the directional pressure on the currency.

Our results confirmed a strong contemporaneous relationship between order flow and exchange rates. For different currency pairs, both similar and different relationships were found. However, the most significant flows tended to come from hedge funds and institutional investors. The relationships between exchange rate moves and flow on previous days were considerably weaker, but our analysis suggested that some flows could be significant for up to 5 days.

In addition a cointegration analysis was conducted (Table 2). This analysis looks for a relationship between the cumulative net flows and the exchange rate over a number of days. Even if a day-by-day relationship were absent any long-term relationship would be highlighted by this test. Although no relationship was found for Dollar/Yen, all the other currencies showed a cointegrating relationship with leveraged investor net flows. In addition Sterling–Dollar showed a relationship with corporate net flows.

The results of the analysis of the order book were similar to the flows analysis. However, this analysis was considered preliminary as only six months of order book data were available and only two currencies were investigated: Euro–Dollar and Sterling–Dollar.

Twelve indicators were derived from the daily ‘snapshot’ of the order book. Two sets of indicators were created by looking at either the *whole* order book or just the *new* orders received during the last 24 h. This

Table 2. Johanson cointegration tests in a VECM specification for order flows (see, e.g. Hayashi (2000), chapter 10). This table shows the number of significant cointegrating equations when the indicated flows are tested against their corresponding FX rates.

Flow series	Number of cointegrating equations under assumption of non-zero intercept but no trend in the FX rates	Number of cointegrating equations under assumption of a non-zero intercept and a linear trend in FX rates
EUR/USD		
Corporates	0	0
Hedge funds	1	0
Asset managers	0	0
Other	0	0
EUR/GBP		
Corporates	0	0
Hedge funds	1	1
Asset managers	0	0
Other	0	0
GBP/USD		
Corporates	1	0
Hedge funds	1	1
Asset managers	0	0
Other	0	0
USD/JPY		
Corporates	0	0
Hedge funds	0	0
Asset managers	0	0
Other	0	0

allowed us to see if new orders carry more information than old ones. Each of these sets was further divided into either only the *take profit* orders or the total orders (both *stop loss* and *take profit*). Finally for each of these four sets we calculated the net order value in three categories: orders with price within 0.5% of the current (mid-point) spot price, orders with price between 0.5% and 1% of the current spot, and the sum of these two (orders within 1% of the spot).

Again, although there was variation between the currency pairs, there were significant correlations between market moves and the order book indicators. Both stop-loss and take-profit indicators were linearly related to future exchange rate moves for lags of up to 5 days. The results here were stronger than those found for the order flow data.

Overall, the statistical results suggest that there is useful structure to be found between the proprietary data sets of flow and limit orders and their corresponding exchange rates. The results did vary a little across the currency pairs, however, and it seemed plausible that the linear approach might be improved upon. For example, only recently placed orders in the Euro–Dollar order book seemed to significantly influence the market whereas for the other currency pair older orders were also important.

The results were nevertheless encouraging and suggested that more advanced techniques might prove fruitful.

6. ERL machine trading with technical indicators, orders and flows

Now an attempt is made to combine transaction flow data and order book data with technical indicators and to use ERL techniques to develop optimal FX trading rules using this information. As the order and flow data is currently only available daily, we used daily FX market data and technical indicators calculated on a daily basis. As for the statistical analysis reported above we have used three markets for the flows (Euro–Dollar, Sterling–Dollar and Yen–Dollar) and two for the order book (Euro–Dollar and Sterling–Dollar). This part of the research is exploratory as the data was restricted to the period March 2002 to August 2002.

Above, we reported the results of using 15 min data with only technical indicators. As we are now using daily data we repeated the ERL optimization of the trading rules for technical indicators on their own. None of the three currencies were profitable at costs above 2 bp per transaction in the case that only technical indicator inputs are considered.

Using daily flow data on its own we find a large improvement in the results (table 3). For Euro–Dollar the out of sample trading is profitable with costs up to 10 bp per transaction (20 bp in total). Sterling–Dollar is just profitable at 20 bp total but Yen–Dollar only up to 4bp total. Interestingly, for Euro, the most significant flows were those of leveraged investors (hedge funds) both net and gross flows, and the net flow of institu-

Table 3. Out-of-sample trading returns using order flow indicators alone.

Currency	Slippage (bp)	Monthly return (%)	Average number of trades per month
GBPUSD	0	1.61	17.2
GBPUSD	2	-2.82	14.4
GBPUSD	4	0.84	6.4
GBPUSD	8	0.58	6.4
GBPUSD	10	0.45	6.4
EURUSD	0	0.88	6.0
EURUSD	2	2.02	4.4
EURUSD	4	1.91	4.4
EURUSD	8	1.69	4.4
EURUSD	10	1.57	4.4
USDJPY	0	0.3	10.8
USDJPY	2	1	12.0
USDJPY	4	-0.3	11.6
USDJPY	8	-0.9	8.0
USDJPY	10	-1.08	8.0

Table 4. Out-of-sample trading returns using order book indicators alone.

Currency	Slippage (bp)	Monthly return (%)	Average number of trades per month
GBPUSD	0	3.14	8.8
GBPUSD	2	0.6	6.4
GBPUSD	4	1.54	5.6
GBPUSD	8	2.38	2.4
GBPUSD	10	0.87	5.6
EURUSD	0	1.82	8.8
EURUSD	2	-0.19	11.6
EURUSD	4	-0.17	10.8
EURUSD	8	1.19	7.2
EURUSD	10	-0.5	8.0

tional investors. For Sterling the institutional flows were the most important followed by corporate flows.

Next order book indicators are considered in isolation (table 4). For Sterling–Dollar the trading was profitable at all costs up to 20 bp total (10 bp per transaction) outperforming both the technicals and the order flows. For Euro–Dollar however the system was only profitable below 4 bp per transaction.

The final tests involved three combinations of indicators: first technical indicators and flows, second technical indicators and the order book and finally flows and the order book. These were investigated for Sterling–Dollar and Euro–Dollar (tables 5–7). The combination of the public information contained in the technical indicators together with the private information contained in either the flows or orders produces a profitable automated trading system, even at 10 basis points per transaction. The results are better than any of the tests with technical indicators, flows or order book

Table 5. Out-of-sample trading returns using technicals and order flow indicators combined.

Currency	Slippage (bp)	Monthly return (%)	Average number of trades per month
GBPUSD	0	1.5	6.0
GBPUSD	2	1.1	3.6
GBPUSD	4	1.05	5.6
GBPUSD	8	1.42	2.4
GBPUSD	10	0.53	2.4
EURUSD	0	1.3	10.4
EURUSD	2	0.81	12.8
EURUSD	4	0.55	12.8
EURUSD	8	0.8	10.4
EURUSD	10	2.93	1.6

Table 6. Out-of-sample trading returns using technicals and order book indicators combined.

Currency	Slippage (bp)	Monthly return (%)	Average number of trades per month
GBPUSD	0	1.99	5.6
GBPUSD	2	1.34	8.8
GBPUSD	4	2.14	5.6
GBPUSD	8	2.38	2.4
GBPUSD	10	2.34	3.2
EURUSD	0	-0.26	6.0
EURUSD	2	2.97	8.0
EURUSD	4	0.41	11.2
EURUSD	8	1.08	6.0
EURUSD	10	-1.76	11.6

Table 7. Out-of-sample trading returns using order flow indicators and order book indicators combined.

Currency	Slippage (bp)	Monthly return (%)	Average number of trades per month
GBPUSD	0	1.75	12.4
GBPUSD	2	0	0
GBPUSD	4	3.51	9.6
GBPUSD	8	1.57	4.8
GBPUSD	10	1.62	4.0
EURUSD	0	0.13	12.0
EURUSD	2	1.82	13.2
EURUSD	4	1.6	2.4
EURUSD	8	1.42	6.8
EURUSD	10	2.47	2.4

separately in isolation. Although only a few technical indicators were chosen by the machine learning system to supplement the orders and flows, these technical indicators prove to be an important addition to other information. It is interesting that using the combined order data and flow data in the absence of technical indicators produces a relatively less profitable trading strategy. Figure 5 compares the profits of the different strategies obtained by using different inputs.




Figure 5. Summary graph of out-of-sample trading returns.

7. Just scratching the surface

The initial studies reported here have confirmed the important role that non-public information, specifically transaction flows and limit orders have in improving forecasts in the FX market. When allied to traditional technical indicators a significant improvement in forecasting performance results from the use of customer flows giving profitable trading rules at up to 20 basis points total costs. The machine trading results are consistent over different currency pairs and seem to be stable to changes in the inputs and the out-of-sample period used for testing.

There are a number of areas in which the collaboration between HSBC and the Centre for Financial Research in Cambridge can be taken forward. These include replication of the flow and order book results using longer time spans of data; a more detailed analysis of the order book using more finely grained indicators rather than the binary indicators used so far; splitting the orders by customer type; and the application of improved machine learning algorithms in the construction of optimal trading rules.

The use of technical indicators for FX trading has a long history and has on various occasions been the subject of academic study. However, it has not been until the inclusion of information that goes beyond the analysis of historic price action that automated trading rules have shown consistent profitability. The information contained in customer transaction flows of a major FX market maker provide enhanced prediction ability, but it is the breakdown into customer type (an inherently

proprietary form of information) that provides the main performance improvement. The additional use of the proprietary customer order book has proved to be the most important component in our system. Our results are also encouraging because they intuitively make sense. Successful traders in the FX markets apply human judgement to a range of information and techniques. In our work we have effectively mimicked these traders by combining the techniques of technical analysis with the stream of information available to them. So far we have only scratched the surface of what performance improvements such information can offer an automated trading system. We hope to take this work further in the future.

In addition in the Centre for Financial Research we are conducting research into the detailed FX market mechanisms. The effect of individual trade size and trade frequency on FX prices is being investigated both empirically and in models that will help us better understand FX market dynamics. The results of this ongoing work will be vital to creating fully adaptive automatic trading systems for the world's deepest financial markets.

References

- Allen F and Karjalainen R 1999 Using genetic algorithms to find technical trading rules *J. Financial Economics* **51** 245–71
- Bates R G, Dempster M A H and Romahi Y 2003 Evolutionary reinforcement learning in FX order book and order flow analysis *Proceedings 2003 IEEE*

- International Conference on Computational Intelligence for Financial Engineering, Hong Kong, March 2003*
pp 355–62
- Dempster M A H and Jones C M 2001 A real-time adaptive trading system using genetic programming *Quantitative Finance* **1** 397–413
- Dempster M A H and Jones C M 2002 Can channel pattern trading be profitably automated? *Eur. J. Finance* **8** 275–301
- Dempster M A H, Payne T W, Romahi Y S and Thompson G W P 2001 Computational learning techniques for intraday FX trading using popular technical indicators *IEEE Trans. Neural Networks*, Special Issue on Computational Finance **12** 744–54
- Dempster M A H and Romahi Y S 2002 Intraday FX trading: an evolutionary reinforcement learning approach *IDEAL 2002 (Lecture Notes in Computer Science 2412)*, ed H Yin (Berlin: Springer) pp 347–58
- Hayashi F 2000 *Econometrics* (Princeton: Princeton University Press)
- Jones C M 1999 Automated technical foreign exchange trading with high frequency data *PhD thesis* Centre for Financial Research, Judge Institute of Management, University of Cambridge
- LeBaron B 1999 Technical trading rule profitability and foreign exchange intervention *J. Int. Econ.* **49** 124–43
- Leemans V 2003 Real time trading systems *MPhil dissertation* Centre for Financial Research, Judge Institute of Management, University of Cambridge
- Lyons R K 2001 *The Microstructure Approach to Exchange Rates* (Cambridge, MA: MIT Press)
- Meese R and Rogoff K 1983 Empirical exchange rate models of the seventies *J. Int. Econ.* **14** 3–24
- Meese R and Rogoff K 1997 The out-of-sample failure of empirical exchange rate models *Exchange Rates and International Macroeconomics* ed J Frenkel (Chicago: University of Chicago Press) pp 23–38
- Neely C and Weller P 2003 Intraday technical trading in the foreign exchange market *J. Int. Money & Finance* **22** 223–37
- Romahi Y 2003 Computational learning techniques in high frequency foreign exchange trading *PhD thesis* Centre for Financial Research, Judge Institute of Management, University of Cambridge