

*Research Papers
in Management Studies*



ADAPTIVE SYSTEMS FOR
FOREIGN EXCHANGE TRADING

M P Austin, R G Bates,
M A H Dempster & S N Williams

WP 15/2003

These papers are produced by the Judge Institute of Management, University of Cambridge.

The papers are circulated for discussion purposes only. Their contents should be considered preliminary and are not to be quoted without the authors' permission.

ADAPTIVE SYSTEMS FOR FOREIGN EXCHANGE TRADING

**M P Austin, R G Bates,
M A H Dempster & S N Williams**

WP 15/2003

To appear in *Eclectic*, September 2004, Institute for Financial Services, London

Dr M P Austin
HSBC Global Markets
8 Canada Square
London E14 5HQ
Tel +44 (0) 207 9915968
Fax: +44 (0) 207 9924980
Email: mark.austin@hsbcgroup.com

R G Bates
Centre for Financial Research
Judge Institute of Management
University of Cambridge
and Cambridge Systems Associates Ltd
1 Earl Street, Cambridge CB1 1JR
Email: r.bates@jims.cam.ac.uk

Professor M A H Dempster
Centre for Financial Research
Judge Institute of Management
University of Cambridge
and Cambridge Systems Associates Ltd
1 Earl Street, Cambridge CB1 1JR
Email: m.dempster@jims.cam.ac.uk

Dr S N Williams
HSBC Global Markets
8 Canada Square
London E14 5HQ
Tel +44 (0) 207 9915967
Fax: +44 (0) 207 9924980
Email: stacy.williams@hsbcgroup.com

Please address enquiries about the series to:

Research Support Manager
Judge Institute of Management
Trumpington Street
Cambridge CB2 1AG, UK
Tel: 01223 760546 Fax: 01223 339701
E-mail: research-support@jims.cam.ac.uk

Adaptive Systems for Foreign Exchange Trading

M. Austin¹, R. G. Bates², M. A. H. Dempster², S. Williams¹

(¹HSBC Global Markets, 8 Canada Square, London E14 5HQ; ²Center for Financial Research, Judge Institute of Management, University of Cambridge, Cambridge CB2 1AG & Cambridge Systems Associates Ltd., 1 Earl Street, Cambridge CB1 1JR)

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 & Rogoff 1983 and 1997). Our interest has been in the techniques used by FX traders and market makers and has focused 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 & Karjalainen 1999 and Neely & Weller 1999.

The work of the Centre for Financial Research in this area is summarised in various publications: Bates RG, Dempster MAH & Romahi Y (2003), Dempster MAH & Jones CM (2001 & 2002), Dempster MAH, Payne TW, Romahi YS & Thompson GWP (2001), Dempster MAH & Romahi YS (2002) and Jones CM (1999).

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: Leemans (2003) and Romahi (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.

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, that is, a trading system that gives an optimal trade-off between risk and return. It is important to remember that we have been looking for ‘understandable’ trading rules and we have been careful to create indicators based on rules commonly adopted in the market. By this we mean something of the form: 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. Making sensible indicators is relatively straightforward though, compared with the task of identifying genuinely useful combinations of them. Finding successful combinations is a very challenging problem 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 we have made significant advances. 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 above papers. 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.

Our results, using only technical indicators, have shown consistent profitability. Using FX market data from 1994 to 1998 at 15 minute 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 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, again at 15 minute frequency, similar results were obtained using technical indicators only. However the trading rules were profitable out of sample only with costs less than 2

basis points per opening or closing transaction. The foreign exchange market has become more efficient over the last few years and the Euro market in particular, being the largest and most liquid market, is highly efficient in its pricing. Although a combination of technical indicators may have been profitable in the past we must now ask if other 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 just technical indicators are encouraging we now consider other sources of information for inclusion into the trading rules and, in particular investigate if non-public information available only to an FX market maker could be used.

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 customers: long-term investors, corporates and short-term speculators such as hedge funds. The joint project between the CFR and HSBC sought to find out if this information could be used to help predict the direction of various FX markets. We looked at this information both on its own and also 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.

To see why this flow and order book data might help in forecasting markets, consider what information each can give us. Flows show what different types of clients are doing. So for example if speculative institutional traders, such as Hedge Funds, are buying Euros and selling Dollars we might reach a very different conclusion about the future direction of Euro-Dollar than if the buying of Euros came from an importer of French wine who had bills to pay in Euros and had little choice in the timing of his purchase of Euros. As we will see transactions from certain types of customer provide more useful information about the future direction of the market than transactions from other types of customer.

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 our initial analysis we used six months of order book information, from March to August 2002. Note that the order book gives a snapshot of potential pressures on the market (buying or selling), it consists of transactions waiting to happen when the market price reaches a particular level.

In the order book we distinguished two types of order: *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 will make the difference clear. We will give the explanation in terms of the Dollar/Yen exchange rate, that is the price of a dollar in terms of Yen.

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 of the flow and order data to give us an indication of how useful this proprietary data might be. Measurable linear structure would be an encouraging sign for the more complex machine learning techniques to come. Statistical analysis of the contemporaneous relationship between FX order flow and price has recently been performed by Lyons (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 we calculated two flow numbers for each of these categories: a gross and a net flow. The gross flow is the total buy transaction volume plus total sell volume. This 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. This gives the directional pressure on the currency.

Our results confirmed a strong contemporaneous relationship between order flow and exchange rates and although it varied a little across the currency pairs, the most significant flows tended to come from hedge funds and institutional investors. The relationship between exchange rate moves and flow on previous days was considerably weaker, but our analysis suggested that some flows could be significant for up to 5 days.

In addition a cointegration analysis was conducted. This 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 is 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 analysis of the order book was similar to the flows. Note, this analysis should be considered preliminary as only six months of order book data was 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 hours. This allowed us to see if new orders carry more information than old ones. Each of these sets were further divided into either just 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 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 across the currency pairs, there were significant correlations between market moves and the orderbook 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 were even stronger than those found for the order flow data.

Overall, the statistical results suggested that there was useful structure to be found between the propriety 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 pairs older orders were also important.

The results were nevertheless encouraging and suggested that a more complex approach might prove fruitful.

6. ERL Machine Trading With Technical Indicators, Orders and Flows

We now look at combining transaction flow data and order book data with technical indicators and use ERL techniques to develop optimal FX trading rules using this information. As the order and flow data is only available daily, we used daily FX market data and technical indicators calculated daily. As for the statistical analysis reported above we 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 minute data with only technical indicators. As we are now using daily data we repeated the ERL optimisation of the trading rules for technical indicators on their own. None of the three currencies were profitable at costs above 2bp per transaction.

Using daily flow data on its own we find a large improvement in the results. For Euro-Dollar the out of sample trading is profitable with costs up to 10bp per transaction (20bp in total). Sterling-Dollar is just profitable at 20bp 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 institutional investors. For Sterling the institutional flows were the most important followed by corporate flows.

We next looked at using order book indicators on their own. For Sterling-Dollar the trading was profitable at all costs up to 20bp total (10bp per transaction) outperforming both the technicals and the order flows. For Euro-Dollar however the system was only profitable below 4bp per transaction.

Our 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. The combination of the public information contained in the technical analysis indicators with the private information in either the flows or orders produces profitable trading, even at 10 basis points per transaction and the results are better than any of the tests of technical indicators alone, flows alone or order book alone. Although only a few technical indicators were chosen by the machine learning system to supplement the orders and flows, these technical indicators are an important addition to the other information. It is interesting that using the combined orders and flows in the absence of technical indicators produces less profitable trading than when the technical indicators are included.

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 point 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 worlds deepest financial markets.

References

- Allen F & Karjalainen R (1999), Using genetic algorithms to find technical trading rules, *J. Financial Economics* **51** 245-271
- Bates RG, Dempster MAH & Romahi Y (2003), Evolutionary Reinforcement Learning in FX Order Book and Order Flow Analysis. In: Proc. 2003 IEEE Int. Conf. on Computational Intelligence for Financial Engineering, Hong Kong March 2003, 355-362
- Dempster MAH & Jones CM (2001), A real-time adaptive trading system using genetic programming, *Quantitative Finance* **1** 397-413
- Dempster MAH & Jones CM (2002), Can channel pattern trading be profitably automated? *European J Finance* **8** 275-301
- Dempster MAH, Payne TW, Romahi YS & Thompson GWP (2001), Computational learning techniques for intraday FX trading using popular technical indicators. *IEEE Transactions on Neural Networks, Special Issue on Computational Finance* **12**(4) 744-754
- Dempster MAH & Romahi YS (2002), Intraday FX trading: An evolutionary reinforcement learning approach. In *IDEAL 2002, Lecture Notes in Computer Science* 2412, ed. H Yin. Springer-Verlag, Berlin, 347-358
- Jones CM (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 International Economics* **49** 124-143
- Leemans V (2003), Real time trading systems. MPhil dissertation, Centre for Financial Research, Judge Institute of Management, University of Cambridge
- R.K.Lyons (2001), *The Microstructure Approach to Exchange Rates*, MIT Press
- Meese R & Rogoff K (1983), Empirical exchange rate models of the seventies. *J International economics* **14** 3-24
- Meese R & Rogoff K (1997), The out-of-sample failure of empirical exchange rate models. In *Exchange Rates and International Macroeconomics*, ed. J Frenkel. Univ of Chicago Press, 23-38
- Neely C & Weller P (2003), Intraday technical trading in the foreign exchange market. *J International Money & Finance* **22** 223-237
- 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

Judge Institute of Management 2003 Working Paper Series

NUMBER	TITLE AND AUTHOR
1/2003	Applying a CBA Model to the 3 Gorges Project in China (R Morimoto & C Hope)
2/2003	Global Strategic Alliance Membership: Market versus institutional criteria (M Warner, E Yin & C J Choi)
3/2003	Culture and Management in China (J Child & M Warner)
4/2003	Nature, Nurture & Economic Growth (B Dowling)
5/2003	Portfolio Management for Pension Funds (S Arbeleche, M A H Dempster, E A Medova & G W P Thompson)
6/2003	Evolutionary Reinforcement Learning in FX Order Book and Order Flow Analysis (R G Bates, M A H Dempster & Y S Romahi)
7/2003	Structured Products for Pension Funds (M A H Dempster, M Germano, E A Medova & M Villaverde)
8/2003	Is the Long Wave Getting Shorter? (B Dowling)
9/2003	A Framework to Measure Integrated Risk (E A Medova & R G Smith)
10/2003	The Marginal Impacts of CO ₂ , CH ₄ and SF ₆ Emissions (C Hope)
11/2003	Review: Structuration Theory and Information Systems Research (M Jones & M Karsten)
12/2003	Global Fund Management Using Stochastic Optimization (M Villaverde)
13/2003	Econometric Modelling for Global Asset Liability Management (S Arbeleche & M A H Dempster)
14/2003	Hedging European and Barrier Options Using Stochastic Optimization (M Villaverde)
15/2003	Adaptive Systems for Foreign Exchange Trading (M P Austin, R G Bates, M A H Dempster & S N Williams)

Working Papers Order Form

Name: _____

Address: _____

Email: _____ Phone: _____

Paper(s) required:

Number	Author(s)	Title

Price per publication = £5.50, including postage and packing.

For orders of 5 papers or more, postage and packing will be billed separately.

Credit card payments up to £20.00 will incur an additional standard charge of £1.00.

Credit card payments over £20.00 will incur an additional standard charge of £3.00.

Total payment enclosed: £ _____

Payment method:

Cheque payable to "University of Cambridge"

Draft/Money order

Credit card Card type Visa Mastercard

Name on card _____

Card number _____

Expiry (mm/yy) _____

Signature _____

Date _____

Please send this completed form, together with your payment, to:

Research Support Manager
Judge Institute of Management, University of Cambridge
Trumpington Street, Cambridge CB2 1AG, UK
Tel: + 44 1223 760546 Fax: + 44 1223 339701
E-mail: research-support@jims.cam.ac.uk