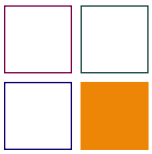

Adaptive systems for foreign exchange trading

A joint project between academics and bankers has shown how banks can improve the forecasting performance of their technical trading systems in foreign exchange markets. **Professor Michael Dempster** and **Graham Bates**, both of the Centre for Financial Research, Cambridge, and **Dr Mark Austin** and **Dr Stacy Williams**, both of HSBC Global Markets, outline the results of their research

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 18 months a joint project with HSBC Global Markets 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.

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 among investment practitioners, technical analysis methods have tended to be regarded with some degree of 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 CFR in this area is summarized in various publications including Dempster MAH and

ExecutiveSummary

- Consistently predicting FX markets has seemed like an impossible goal but recent advances in financial research now suggest it can be done. Automated trading systems are being used successfully to predict intraday and daily exchange rates.

- Trading systems using only publicly available technical indicators can be profitable – but those that also use proprietary information can be more accurate and therefore more profitable.

- A joint project by the Centre for Financial Research (at the Judge Institute of Management, Cambridge University) and HSBC used the bank's customer order information to show that using proprietary information in trading systems can improve their forecasting performance and profitability.

- The research findings also intuitively make sense. Successful traders in the FX markets apply human judgement to a range of information and techniques. In this project the researchers mimicked these traders by combining the techniques of technical analysis with the stream of public and non-public information available to them.

Jones CM (2001 & 2002). Much of it 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 optimizing 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 that looks most promising.

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 exchange rates, trading was profitable on all three currency pairs, out-of-sample, even after allowing for costs of up to 2 basis points (bp) for the opening and closing of transactions (4 bp overall). With some machine learning techniques profitability extends to costs of 8 bp overall in some currencies.

Using more recent data on euro/dollar rates 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 bp per opening or closing trans-

action. 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.

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

Making it happen: HSBC's foreign exchange dealing room

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.

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.

HSBC FX traders at their desks: human judgement was mimicked by the researchers

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).
- Companies (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 weaker, but our analysis suggested that some flows could be significant for up to five days.

In addition, a co-integration 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 co-integrating 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. However, 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 per cent of the current spot price, orders with price between 0.5 per cent

and 1 per cent of the current spot, and the sum of these two (orders within 1 per cent 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 five 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.

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 optimization of the trading rules for technical indicators on their own. None of the three currencies were profitable at costs above 2bp (basis points) 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 10bp 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.

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 20bp 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 areas in which the collaboration between HSBC and the CFR 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 been widely studied. 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 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 some way further in the near 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, which will help us better understand FX market dynamics. The results of this continuing work will be vital to creating fully adaptive automatic trading systems for the world's deepest financial markets.

Dr Mark Austin *is managing director, foreign exchange strategy, at HSBC Global Markets.*

Graham Bates *is researcher at the Centre for Financial Research (CFR), Judge Institute of Management, University of Cambridge, and also an associate at Cambridge Systems Associates Ltd.*

Professor Michael Dempster *is director of CFR and managing director of Cambridge Systems Associates. Ltd.*

Dr Stacy Williams *is quantitative strategist, foreign exchange, at HSBC Global Markets.*

RecommendedReading

- Allen F & Karjalainen R, "Using genetic algorithms to find technical trading rules", *Journal of Financial Economics*, Vol. 51, 245-271, 1999.
- Dempster MAH & Jones CM, "A real-time adaptive trading system using genetic programming", *Quantitative Finance,* Vol 1, 397-413, 2001.
- Dempster MAH & Jones CM, "Can channel pattern trading be profitably automated?", *European Journal of Finance*,Vol 8, 275-301, 2002.
- LeBaron B, "Technical trading rule profitability and foreign exchange intervention", *Journal of International Economics*, Vol 49, 124-143, 1999.
- Leemans V, "Real-time trading systems", MPhil dissertation, CFR, Univ. of Cambridge, 2003.
- Lyons RK, *The Microstructure Approach to Exchange Rates*, MIT Press, 2001.
- Meese R & Rogoff K, "Empirical exchange rate models of the Seventies", *Journal of International Economics*, Vol 14, 3-24, 1983.
- Meese R & Rogoff K, "The out-of-sample failure of empirical exchange rate models", *Exchange Rates and International Macroeconomics*, 23-38, ed. J Frenkel, Univ of Chicago Press, 1997.
- Neely C & Weller P, "Intraday technical trading in the foreign exchange market," *Journal of International Money & Finance*, Vol 22, 223-237, 2003.
- Romahi Y, "Computational learning techniques in high frequency foreign exchange trading," PhD thesis, CFR, University of Cambridge, 2003.