Multi-agent Foreign Exchange Market Modelling via GP

Stephen Dignum and Riccardo Poli

Department of Computer Science
University of Essex

Abstract. In this work genetic programming is used to express and evolve trading strategies for a foreign exchange currency market simulator.

1. Introduction

Multi-agent systems built using agents with simple behaviours can be used to demonstrate complicated, often counter intuitive, results seen in the real world [1]. A currency exchange market offers an ideal environment to put this idea into practice. Here, traders and their clients can be thought of as agents with the desire to exchange currencies to make a profit. A successful trader will pay no heed to past events, but simply react to current conditions quickly to take advantage of potential opportunities [2]. So, our approach is to use GP to evolve agents that react to trade requests from other agents, current trading rates, and news that could affect the value of currency.

2. Our Foreign Exchange Currency Market Simulator

A central location (the Currency Market) records and calculates trading rates between currencies to facilitate currency transactions. When a trade is transacted, a trading agent sends details of the agreed trade to the Currency Market, which updates the rate of the currency that has been purchased. To allow external influences upon the market, a component of the simulator provides current news to agents for each trading round. This information has been abstracted to a numerical form.

A *client agent* chooses a pair of currencies to trade, one to buy, another to sell, based on current market conditions. The choice is stochastic but is biased by current currency news. The client then selects a trader to transact with. This decision is based on previous transactions. If the trader refuses to trade another one is chosen.

Trading agents are more sophisticated. They have holdings in particular currencies and can increment and decrement those holdings by trading with both clients and other traders. Traders follow particular strategies in order to maximise their overall balance. To implement this functionality, trading agents have a mechanism to store their current holdings and to update those holdings when a transaction is successfully completed. The agent is also equipped with the ability to determine the value of its holdings. This value acts as a *fitness function* and determines whether the agent is allowed to continue trading i.e. whether it is still solvent.

On creation each trader is provided with a GP tree representing its trading strategy. This is used in two ways. First when a request to trade is received, if the amount returned from the expression is greater or equal to the amount requested in the trade, the trader accepts that trade. Secondly, when the trader is given the opportunity to trade on its own account, it chooses two currencies at random and selects the appropriate trade amount by evaluating the expression.

After initialisation, the simulator repeatedly executes trading rounds. A *trading round* is a session where all clients and traders have the opportunity to request a trade. A real world analogy for a trading round would be of a trading session within a financial market, the output of each round can be thought of as a set of 'close of business' valuations.

Not all trading rounds involve evolution of the trading expressions. So, an agent's performance can be judged over a number of rounds. In an *evolution round*, bankrupt traders and traders that have not traded by a specified number of rounds are removed from the main population. The simulator then creates new trading agents by using standard GP operators and an appropriate wealth-initialisation strategy.

3. Results and Conclusions

A variety of preliminary experiments have been performed with our GP-based trading simulator (see [3] for a fuller set of results). These have emphasized a number of interesting effects. For example, in many simulations we have seen a large number of occurrences of the sub-tree (*DIVIDE buyRate buyHolding*). This expression reduces an agent's desire to purchase a currency when his holding of the currency is high, which is a good trading strategy.

Modelling trader and client behaviours within a simulated trading environment can provide a valuable technique to analyse the dynamics of foreign exchange markets. Using a multi-agent approach may allow us to engineer and replicate, from simple behaviours, complicated, often emergent, aspects of financial markets.

Evolutionary methods allow agents to adapt to new market conditions, and provide a facility for the researcher to not only identify expressions and variables pertinent to certain market conditions, but also to discover robust trading strategies that are successful over a large range of trading circumstances. This paper has concentrated on GP as the primary technique for evolutionary change.

Bibliography

[1] J. Rauch, Seeing around Corners, *The Atlantic Monthly*, April 2002, pages 35-48 [http://www.theatlantic.com/issues/2002/04/rauch.htm], Last visited: 26/08/2003

[2] M. Lewis (1999), *Liars Poker: Two Cities True Greed*, Coronet Books, Hodder and Stoughton Ltd, 338 Euston Road, London NW1 3BH, 1999

[3] S. Dignum and R. Poli, Multi-agent Foreign Exchange Market Modelling via GP, Department of Computer Science, University of Essex, Tech. Rep. CSM-400, March 2004.