Trading Foreign Exchange Markets Using Evolutionary Automatic Programming

Tony Brabazon
Department of Accountancy
University College Dublin
Ireland
Email: Anthony.brabazon@ucd.ie

Michael O'Neill
Dept. of Computer Science & Information Systems
University of Limerick
Ireland
Email: michael.oneill@ul.ie

Abstract

This study examines the potential of Grammatical Evolution to uncover a series of useful technical trading rules which can be used to trade foreign exchange markets. In this study, each of the evolved programs represents a market trading system and implicitly, a predictive model. The form of these programs is not specified ex-ante but emerges by means of an evolutionary process. Daily US-STG exchange rates for the period 9/3/93 to 13/10/97 are used to train and test the model. The preliminary findings suggest that the developed rules earn positive, risk-adjusted returns in hold-out sample test periods, after allowing for trading and slippage costs.

1 Introduction

The objective of this study is to determine whether Grammatical Evolution is capable of uncovering useful technical trading rules for trading in spot foreign exchange markets. Spot markets involve transactions in which two currencies are immediately exchanged. Foreign exchange markets are the most active of all financial markets with average daily trading volumes in traditional (non-electronic broker) foreign exchange markets estimated at $1.2 trillion, with the the average daily trading volume in spot markets amounting to $387 billion. International foreign-exchange markets are dominated by a small number of heavily-traded currency pairings. Approximately 90% of all foreign currency transactions involve the US Dollar [2]. This study develops trading systems for the US Dollar / Sterling currency pair.

1.1 Technical analysis

Technical analysis is defined as the attempt to identify regularities in the time-series of prices by extracting patterns from noisy data [11]. The price of one currency in terms of another is termed an exchange rate. Some market traders, known as technical analysts, believe that prices (exchange rates) move in trends and that price patterns repeat themselves [13]. If we accept this premise, that there are rules, although not necessarily static rules, underlying price behaviour it follows that opportunities may exist to enhance trading decisions through use of an appropriate rule induction methodology such as Grammatical Evolution (GE). Although controversy exists amongst financial theorists regarding the veracity of the claim of technical analysts, the methods are widely applied in practice. [19], in a study conducted on behalf of the Bank of England, found that approximately 90% of financial institutions dealing in foreign exchange in London placed some weight on information obtained from technical analysis. London is the largest foreign-exchange market in the world, with 31% of total trading volume. The second largest market New York, has approximately 16% of total trading volume [2]. Greatest use was made of this information in forming short-run (a trading horizon of less than three months) exchange rate expectations.

1.2 Motivation for study

As noted by [8] there are a number of reasons to suppose that the use of an evolutionary automatic programming (EAP) approach such as GE or Genetic Programming (GP) can prove fruitful in the financial prediction domain. EAP can conduct an efficient exploration of a search space and can uncover dependencies between input variables, leading to the selection of a good subset for inclusion in the final model. Additionally, use of EAP facilitates the implementation of complex fitness functions including discontinuous, non-differentiable functions. This is of particular importance in a financial domain as fitness criterion may be complex, usually requiring a balancing of risk and return. Another useful feature of an EAP approach is that it can produce human-readable rules that have the potential to enhance understanding of the problem domain.

A number of studies have applied genetic algorithms (GA) or GP to the generation of technical trading rules in equity markets [1] [6] [3] [14] [7]. However, little published work exists concerning the application of these methods to technical trading of foreign exchange markets.

The need to research the potential of EAP methodologies in the foreign exchange markets has been recognised in the recent past [9]. This study introduces Grammatical Evolution to the problem domain.

1.3 Technical Indicators

The development of trading rules based on current and historic market price information has a long history [5]. The process entails the selection of one or more technical indicators and the development of a trading system based on these indicators. These indicators are formed from various combinations of current and historic price information. Although there are potentially an infinite number of such indicators, the financial literature [4] [13] [17] suggests that certain indicators are widely used by investors. Four groupings of indicators are given prominence in the technical analysis literature:

i. Moving average indicators

ii. Momentum indicators

iii. Trading range indicators

iv. Oscillators

We intend to use of each of these groupings as our model is developed, but in our preliminary study described in this paper, we have limited our attention to the first three groupings of indicators.

The rest of this paper is organised as follows. Section two describes the evolutionary algorithm adopted, Grammatical Evolution [18] [16] [15] and section three outlines the data set and methodology utilised. The following sections provide the results of the study followed by a number of conclusions.

2 Grammatical Evolution

A description of Grammatical Evolution (GE) can be found in [18] [16] [15]. In this study, each of the evolved programs represents a market trading system. The model structure / form of this program is not specified ex-ante but emerges by means of an evolutionary process. Unlike GAs, which generally have a fixed string length, implying an ex-ante model specification, neither GP nor GE impose a fixed structure ex-ante. In contrast with the parse tree representation of the program adopted in traditional GP [10], GE utilises a linear genome representation. GE can be further distinguished from GP in that it adopts a genotype-phenotype mapping process to generate the output program for each individual in the population. Each individual, a variable length binary string, contains in its codons (groups of 8 bits) the information to select program production rules from a Backus Naur Form (BNF) grammar.

The grammar adopted in this study is given below.

```
N={<code>,<expr>,<fopbi>,<fopun>,
   <matbi>,<relbi>,<var>,<int>}
T={p,=,(,),f_and,f_or,f_not,+,-,*,>,
   <,>=,<=,scale,ma,day,1,2,3,4,5,10}
S=<code>
P=

<code> ::=  p = <expr> ;
<expr> ::=  <fopbi> ( <expr> , <expr> )
          | <fopun> ( <expr> )
          | <expr> <matbi> <expr>
          | <expr> <relbi> <expr>
          | <var>
<fopbi> ::= f_and
          | f_or
<fopun> ::= f_not
<matbi> ::= + | - | *
<relbi> ::= > | < | >= |<=
<var> ::= scale( <int> )
        | scale(ma( <int> , day ))
        | scale(day)
        | scale(momentum( <int> , day ))
        | scale(trb( <int> , day ))
<int> ::= 1 | 2 | 3 | 4 | 5 | 10
```

In this study, the GE algorithm uses a steady state replacement mechanism, such that, two parents produce two children the best of which replaces the worst individual in the current population, if the child has greater fitness. The standard genetic operators of mutation (probability of 0.01), and crossover (probability of 0.9) are adopted as is a duplication operator (probability of 0.01), that duplicates a random codon and inserts this into the penultimate codon position on the genome.

3 Experimental Approach

This study uses daily closing exchange rate data drawn from the London market for the period 9/3/93 to 13/10/97. The foreign-exchange rates are initially normalised using a 75-day moving average followed by linear scaling into the range 0 to 1. The training data set was comprised of the first 800 trading days of the data set. The remaining data was divided into two hold-out samples totalling 400 days. The division of the hold-out period into two segments was undertaken to allow comparison of the hold-out results across different market conditions in order to assess the stability and degradation characteristics of the developed models' predictions.

The rules evolved by GE are used to generate one of three signals for each day of the training or test periods. The possible signals are *Buy*, *Sell*, or *Do-Nothing*. Permitting the model to output a Do-Nothing signal reduces the hard threshold problem associated with production of a binary output.

A variant on the trading methodology developed in [4] is then applied. If a buy signal is indicated on a given day, a fixed investment of $1,000 is made in the foreign currency using borrowed funds. This position is automatically closed at the end of a five day trading period. On the production of a sell signal, an investment of $1,000 is made on the short side and again this position is closed out after a five day period. This gives rise to a maximum potential investment of $5,000 at any point in time. The potential loss on an unprotected short sale is in theory infinite but in practice is unlikely to exceed

$1,000. On a short sale, an amount of $1,000 is considered invested for the five day period of the transaction. The total return generated by the developed trading system is a combination of its trading return net of trading costs and an interest differential.

The technical indicators adopted in this study are the moving average, momentum and trading range breakout where the relevant periods and combination of indicators used are to be determined by evolution. Apart from the raw components of the technical indicators adopted (average, moving average, lag, each of which takes a real-valued argument representing the chosen time-window) the BNF grammar also allows the use of standard arithmetic operators (+, -, * and inequality operators) and the binary operators f_and, f_or, and the unary operator f_not. The operations f_and(argument 1, argument 2), f_or (argument 1, argument 2), and f_not (argument 1) return the minimum, maximum, of the arguments, and 1 - the argument, respectively. The current day's price is also provided as a model input.

The signals generated for each day are post-processed using the following rule:

$$Sell = 0.0 >= Value < .33$$

$$Do - Nothing = .33 >= Value < .66$$

$$Buy = .66 >= Value <= 1.0$$

3.1 Selection of Fitness Function

The quality of each developed trading system is measured against a fitness criterion. A simple fitness metric such as profitability, both in and out of sample, is inadequate as it fails to consider the risk of the developed trading system. This section develops both the return and risk metrics incorporated in the fitness function used in this study.

Consider a daily observed exchange rate series S_t with $t = 1, 2, ...n$. We define the five day return for a long position as $(\frac{S_t}{S_{t+5}} - 1)$ and as $-1(\frac{S_t}{S_{t+5}} - 1)$ for a short position. Therefore, the total return earned from an individual long trade is calculated as:

$$R_{(t,t+5)} = (\frac{S_t}{S_{t+5}} - 1) - (i_{home} - i_{foreign})$$

and for a short trade, the total return is calculated as:

$$R_{(t,t+5)} = -1(\frac{S_t}{S_{t+5}} - 1) - (i_{home} - i_{foreign})$$

where, $R_{(t,t+5)}$ is the return, i_{home} represents the dollar interest rate and $i_{foreign}$ represents the foreign interest rate. In this trading strategy, the home currency is the US dollar. When a long position is taken, the foreign currency is purchased and the return equals the trading return plus the interest earned in the foreign currency, less the cost of the home funds borrowed to invest in the foreign currency. When a short signal is produced by the trading system, the foreign currency is borrowed, at a cost of the foreign interest rate, and is invested in US dollars for the five day trading period. The dollar interest rate is earned on these funds for five days. i_{home} represents the dollar interest rate and $i_{foreign}$ represents the foreign interest rate. A trading cost of 0.025% in each direction and a slippage allowance of 0.01% is included [12].

This study adopts a conservative approach to risk and seeks to minimise the risk of a catastrophic loss by incorporating a measure of the system's drawdown (maximum cumulative loss since commencing trading with the system) into the fitness function:

$$Fitness = ExcessReturn - x * (Max.drawdown)$$

'x' can be considered as a tuning factor. The more risk-adverse the investor, in terms of wishing to avoid a catastrophic loss, the higher the value of this parameter. We initially set this parameter to 1. The chosen risk metric will help ensure a stringent test of trading rule performance as high reward / risk trading rules will be discriminated against during the evolutionary process.

4 Results

The results from our preliminary experiments are now provided. 30 independent trials were conducted with a population size of 500 individuals and an evolutionary period of 50 generations. To assess the quality of the results obtained they are compared against a 'buy-and-hold' benchmark. The 'buy-and-hold' benchmark is determined by comparing the return obtained from investing $5,000 in the foreign currency, net of the interest rate differential, during each training and out-of-sample period. A plot of the mean best fitness at each generation over the 30 runs can be seen in Figure 1.

Table 1 provides a comparison of the performance of the best evolved rule set against that of the benchmark investment strategy.

Trading Period	*Evolved Rule set*	*Buy-and-Hold*
Training	11.2628	0.04377
Test 1	2.954636	0.089145
Test 2	3.581367	0.027894

Table 1: Results for the best evolved rule set over the 30 runs compared to the benchmark buy-and-hold strategy.

The evolved trading rules generated returns on the hold-out samples, after allowance for trading, slippage and net interest costs, of. The returns are significantly in excess of the benchmark buy-and-hold strategy over each of the test and training periods.

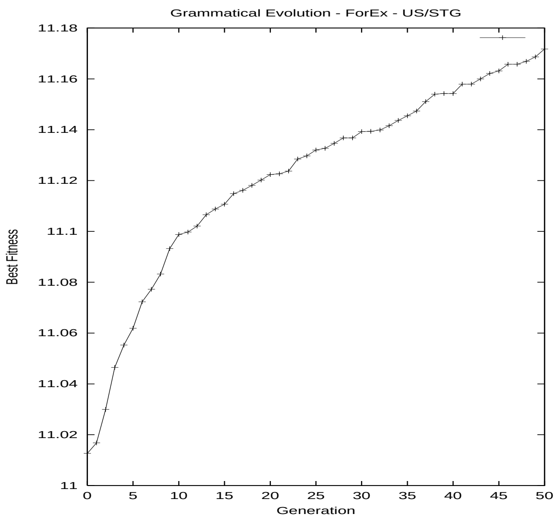

Figure 1: A plot of the mean best fitness over the 30 runs.

5 Conclusions & Future Work

GE was shown to evolve trading rules that were superior to the benchmark buy-and-hold strategy over all the trading periods analysed. There is notable scope for further research utilising GE in this problem domain. Our methodology has included a number of simplifications, for example, we only considered a small set of primitive technical indicators. The incorporation of additional technical indicators may further improve the performance of our approach. In a similar fashion, a predictive model with a longer horizon could be constructed by applying grammatical evolution to a series of fundamental indicators.

References

[1] Allen, F., Karjalainen, R. (1999). 'Using genetic algorithms to find technical trading rules', *Journal of Financial Economics*, 51:245-271.

[2] Bank for International Settlements (2001). 'Central bank survey of foreign exchange and derivatives market activity in April 2001', *Press Release, 31/2001E, Bank of International Settlements*, October 2001.

[3] Bauer R. (1994). *Genetic Algorithms and Investment Strategies*, New York:John Wiley & Sons Inc.

[4] Brock, W., Lakonishok, J. and LeBaron B. (1992). 'Simple Technical Trading Rules and the Stochastic Properties of Stock Returns', *Journal of Finance*, 47(5):1731-1764.

[5] Brown, S., Goetzmann W. and Kumar A. (1998). 'The Dow Theory: William Peter Hamilton's Track Record Reconsidered', *Journal of Finance*, 53(4):1311-1333.

[6] Colin, A. (1994). *Genetic Algorithms for Financial Modelling*, in Guido Deboeck, ed., Trading on the edge: neural, genetic and fuzzy systems for chaotic and financial markets, New York:John Wiley & Sons Inc.

[7] Deboeck G. (1994). *Using GAs to optimise a trading system*', in Guido Deboeck, ed., Trading on the edge: neural, genetic and fuzzy systems for chaotic and financial markets, New York: John Wiley & Sons Inc.

[8] Iba H. and Nikolaev N. (2000). 'Genetic Programming Polynomial Models of Financial Data Series', In *Proc. of CEC 2000*, 1459-1466, IEEE Press.

[9] Jegadeesh, N. (2000). 'Discussion - Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation', *Journal of Finance*, 55(4):1765-1770.

[10] Koza, J. (1992). *Genetic Programming*, Massachusetts: MIT Press.

[11] Lo, A. W., Mamaysky, H. and Wang, J. (2000). 'Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation', *Journal of Finance*, 55(4):1705-1765.

[12] Levich, R. and Thomas, L. (1993). 'The significance of technical trading-rule profits in the foreign exchange market: a bootstrap approach', *Journal of International Money and Finance*, 12:451-474.

[13] Murphy, John J. (1999). *Technical Analysis of the Financial Markets*, New York: New York Institute of Finance.

[14] Neely, C., Weller P. and Dittmar, R. (1997). 'Is technical analysis in the foreign exchange market profitable? A genetic programming approach', *Journal of Financial and Quantitative Analysis*, 32(4):405-428.

[15] O'Neill M. (2001). Automatic Programming in an Arbitrary Language: Evolving Programs with Grammatical Evolution, PhD thesis, University of Limerick, Ireland, 2001.

[16] O'Neill M., Ryan C. (2001).'Grammatical Evolution', *IEEE Trans. Evolutionary Computation*, 5(4), 2001.

[17] Pring, M. (1991). *Technical analysis explained: the successful investor's guide to spotting investment trends and turning points*, New York:Mc Graw-Hill Inc.

[18] Ryan C., Collins J.J., O'Neill M. (1998). Grammatical Evolution: Evolving Programs for an Arbitrary Language. *Lecture Notes in Computer Science 1391, Proceedings of the First European Workshop on Genetic Programming*, 83-95, Springer-Verlag.

[19] Taylor, M. and Allen, H. (1992). 'The use of technical analysis in the foreign exchange market', *Journal of International Money and Finance*, 11:304-314.