

Interday Foreign Exchange Trading using Linear Genetic Programming

Garnett Wilson
Department of Computer Science
Memorial University of Newfoundland
St. John's, Canada
gwilson@mun.ca

Wolfgang Banzhaf
Department of Computer Science
Memorial University of Newfoundland
St. John's, Canada
banzhaf@mun.ca

ABSTRACT

Foreign exchange (forex) market trading using evolutionary algorithms is an active and controversial area of research. We investigate the use of a linear genetic programming (LGP) system for automated forex trading of four major currency pairs. Fitness functions with varying degrees of conservatism through the incorporation of maximum draw-down are considered. The use of the fitness types in the LGP system for different currency value trends are examined in terms of performance over time, underlying trading strategies, and overall profitability. An analysis of trade profitability shows that the LGP system is very accurate at both buying to achieve profit and selling to prevent loss, with moderate levels of trading activity.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search—*Heuristic methods*

General Terms

Algorithms, Performance, Experimentation

Keywords

foreign exchange markets, computational finance, linear genetic programming, technical analysis

1. INTRODUCTION

The foreign exchange (“forex” or “FX”) trading market is different from stock market trading in a few ways that give its price data unique characteristics: FX traders are able to conduct trades 24 hours a day using exchanges all over the world, FX traders can buy or sell currencies at the current rate almost instantly (no time for agreement on the trade required), and forex markets represent the most active of all financial markets with average daily trading volumes exceeding 1 trillion dollars and over 90% of transactions involving

the US dollar (USD) [1]. As a result of these characteristics, there are no abrupt interday price changes or lags in price change due to lack of available parties for a trade. Furthermore, the value of one currency increases or decreases in relation to another. This paper describes (to the authors’ knowledge) the first application of linear genetic programming (LGP) to the field of forex trading, examining the performance of the system on the 2008 exchange rates of four major currency pairs. The results indicate a profitable system that is very accurate at making trading decisions in forex markets.

The motivation for the use of LGP applied to foreign exchange markets is to analyze the behavior of an LGP trading system that has previously shown promise in stock market trading, as described in [16]. It is often recommended to test the robustness of trading systems using alternative markets [7]. Such markets can provide exposure to new conditions and trends to help determine strengths (and weaknesses) of trading systems, especially when they are not correlated with other markets on which the system is being tested. The purpose of this paper is thus to explore the real world application of LGP to an untried financial market domain. To this end, we examine previously proposed fitness function modifications particular to the foreign exchange market domain and apply them to the LGP trading system.

The next section describes existing literature on the application of evolutionary algorithms to forex trading. Section 3 describes the linear GP system applied to forex trading, with Section 4 detailing considerations for fitness evaluation. Section 5 describes results; conclusions follow in Section 6.

2. GENETIC PROGRAMMING FOR FOREX TRADING

Evolution-inspired algorithms have been applied to automated trading solutions in markets since the 1990s, with a recent resurgence in interest seen in the literature of the past few years. Neely et al. [11] created a tree-based genetic programming (GP) implementation for trading long and short positions in a foreign currency with base currency in US dollars (USD) in 1997. Using forex data for a period of 21 years (early 1974 to late 1995), the authors examined the interday values (while holding overnight) of the DEM, JPY, GBP, and CHF against USD and two cross-rate series, DEM/JPY and GBP/CHF. The authors achieved mean annual excess returns of 1.01% to 6.05%. In later work, Neely et al. [10] used a slightly modified version of their GP algorithm to examine half-hour bid and ask quotes for spot foreign exchange rates during 1996. Their results for in-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO’10, July 7–11, 2010, Portland, Oregon, USA.

Copyright 2010 ACM 978-1-4503-0072-8/10/07 ...\$10.00.

trading data, unlike the previous interday results, indicated that the GP did not generate positive excess returns when factoring in transaction costs. In a very recent work that examines a number of approaches utilizing technical trading rules [12], the authors claim that profitability of particular types of trading rules (filter and moving average-based) disappeared during the 1990s, but their returns in [11] (up to 1995) were genuine.

Dacorogna et al. [3] describe foreign exchange trading strategies based on research of the well known FOREX trading firm OANDA using data from the beginning of 1987 to the end of 1995. In [3], the authors describe a tree-based GP system that generated yearly out-of-sample returns from 3.10% to 6.29% using intraday (hourly) forex data for DEM/USD, JPY/USD, and CHF/USD. Brabazon and O'Neill [1] use a special type of genetic programming that evolves trading rules based on a formal grammar called Grammatical Evolution (GE). They describe the application of GE to forex data from an overlapping time period (interday closing exchange rate from late 1992 to late 1997) using US dollar to DEM, JPY, and GBP. The authors found that they were able to generate returns of 0.1% to 5.0% over their test period of 548 days depending on currency. Dempster and Jones [4] used a GA-based system that evaluated technical indicators every 15 minutes while trading at one minute intervals. Different optimization variants in their system were tried, and the best strategy was found to return 7% per annum on the GBP/USD rate from 1993 to 1997.

A number of papers have since claimed successful application of evolutionary computation to forex trading beyond 1997. Schwaerzel and Bylander [15] use traditional GP to analyze GBP/USD and JPY/USD daily rates from the start of 1990 to late 2005. Profits of 89% for GBP/USD and 80% for JPY/USD were achieved for the 10 year test period, but the authors state that these returns are low for a such a long test period. Hryshko and Downs [9] applied a GA to the evolution of entry and exit trading rules, along with reinforcement learning, to the EUR/USD rate from June 2 to December 31, 2002 with a five minute frequency. The authors found that the system achieved profitability of 7% for the first 2 months out of sample, and 6.5% in the 3.5 months out of sample. Dempster et al. examined a GA and reinforcement learning-based system in [5] for the foreign exchange market after 1997. They found that the system tended to over-fit solutions when examining data for GBP/USD from 1994 to 1998. The profitability of the results varied considerably based on slippage used. In a subsequent work, Dempster et al. [6] implemented a system that attempts to correct their previous issues with overfitting. The system produced was a hybrid GA and reinforcement learning system where the GA evolved a set of indicators provided to the reinforcement learning component of the algorithm. The implementation was tested on GBP/USD, USD/CHF, and USD/JPY rates from 1994 to 1998 (same time frame as their previous study in [5]) and produced profitable strategies of approximately 5% to 15% given higher slippage levels.

Recently, Saks and Maringer [14] examined intraday bid-ask tick data for the calendar year of 2006, with a sample approximately every minute for USD/EUR. Their implementation took a novel approach involving evolution of trees representing combinations of long and short exit and entry strategies. The authors' algorithm outperformed randomly

initialized strategies and typically produced average profits from -5 to 15%. Hirabayashi et al. [8] analyzed the hourly closing price of USD/JPY and EUR/JPY with data from 2005-2008. Their GA system outperformed a neural network, with a GA using leveraged trades performing better than a non-leveraged version. Profits using a GA with leveraging ranged from less than 10% to over 70% at particular points in the data series.

3. LGP FOREX TRADING SYSTEM

The algorithm applied to foreign exchange trading in this work is linear genetic programming (LGP). LGP evolves individuals composed of binary strings and associated registers, in contrast to the tree-based individuals of traditional GP. The bits that make up an individual can be further divided into separate instructions. These instructions can be executed sequentially, with the subresult of each instruction stored in one of the individual's registers. Each instruction performs an operation corresponding to a member of the function set that it references, and uses subresults in registers or external data as operands. Thus, each LGP individual represents a small machine-language type program that is evaluated by a fitness function upon termination, often by examining a final solution in one of the individual's registers. LGP is now considered an established hallmark form of GP, and further details of general LGP implementation can be found in [2].

Linear genetic programming is used by its proponents for a number of reasons. Firstly, the internal registers used in LGP individuals provide a way of storing sub-results and thus reduce the need to introduce new operators into a function set for the purposes of solving a problem. Considerable flexibility is also provided by LGP, since each instruction in an LGP individual does not necessarily contribute to the final result in the individual's solution register. Due to this trait of LGP individuals, redundant and detrimental code can be present in the genome of an individual without affecting the individual's fitness. In contrast, each node in a tree-based GP individual contributes to the value found in the root node as a solution.

In this work, we apply a LGP implementation to four forex time series in interday frequencies: CAD/USD, EUR/USD, GBP/USD, and JPY/USD. An initial period was reserved to establish values of technical indicators, following which the GP fitness was evaluated on data corresponding to a moving window of 5 days. Individuals represented sets of trading rules, based on functions in the function set (to be described). For each window of 5 trading days m to n , each of m to $n - 1$ days were used for calculation of a trading decision, with $m + 1$ to n left for the evaluation of the recommendation based on the preceding days. Data used for the determination of a trading decision were normalized using two-phase preprocessing similar to treatment of stock data in [1]: All daily values were transformed by division using a lagged moving average, and then normalized using linear scaling into the range $[0, 1]$ using

$$v_{scaled} = \frac{v_t - l_n}{h_n - l_n} \quad (1)$$

where v_{scaled} is the normalized trading value, v_t is the transformed trading value at time step t , h_n is highest transformed value in the last n time steps, l_n is the lowest trans-

08/01/2008,0.9985	Training Window ↓ Unknown Value
09/01/2008,1.0083	
10/01/2008,1.0139	
11/01/2008,1.0198	
14/01/2008,1.021	
15/01/2008,1.0167	

Figure 1: The LGP trains repeatedly on a window of five values with four training cases (four boxes within the larger grey box). Unknown value is the price for which the “live” trading system receives a recommendation from the LGP.

formed value in the last n time steps, and n is the length of the time lag chosen for the initial transformation.

In addition to an instruction set, each LGP individual possesses four registers, a flag for storing the current value of logical operators, and an output (trade) register for storing the value corresponding to a trade recommendation following execution of the instruction set. If the value of the trade register is 0, no trade is conducted. If the value in the trade register corresponds to a value in the range $+/-[0, 1]$, that value is multiplied by a maximum dollar amount to be bought or sold per trade (\$100 was used based on a starting balance of \$1000 with which to trade) to give some portion of \$100 to be traded. For each trade conducted, there is a transaction cost of 0.025% and slippage of 0.01%, as recommended in [1]. In addition, interest is charged for USD dollars borrowed and interest is earned in the foreign currency while it is held.

The best individual, consisting of the best trading rule set, is used by a “live” trading algorithm. The live trading system provides known information to the LGP for days m to n . The LGP algorithm returns a recommendation for the live trading system, which is used as the basis of its trade on the following day, $n + 1$. The net number of units of foreign currency bought and sold by the best evolved individual (trading rules) given the recommendation of the trade register over all the fitness cases (4 fitness cases per 5 day window) is the buy or sell recommendation to the “live” trading system. The best LGP individual can thus recommend up to \$400 worth of currency selling or buying per actual trading day to the live system. With the next window shift, the current currencies held by the “live” trading system are the new initial amounts for the GP individuals in the next tournament based on the rates in the new window of values. The transactions of the live trading system are actually based on unknown data and determine the success of the forex trader. The process is summarized in Figure 1.

The LGP function set includes standard mathematical operators and logical operators. In addition, established technical analysis metrics such as moving average, momentum, and channel breakout were used. Moving average is the mean of the previous n share prices. The momentum is the ratio of a time-lagged price to the current price. Channel breakout typically uses Bollinger bands around a n -day moving average of the price at $+/- 2$ standard deviations of the price movement over the last n days to alert the trader of significant movements in rates. Each LGP tournament consisted of 1000 rounds. XOR mutation on individual instructions was used with a probability of 0.5, and crossover occurred with a probability of 0.9.

4. FITNESS EVALUATION

The fitness of a GP individual for trading applications is often simply how profitable the trading strategy represented by the individual is over time, or some modification on the measure of that direct profitability. This profitability of the individual / trading rules is naturally determined (at least indirectly) by the established technical indicators that are members of the function set (typically moving average, momentum, and others). Another method of creating profitable individuals (trading rules) is to modify the fitness measure directly. For instance, Saks and Maringer [14] tailor the fitness function to risk tolerance of traders. Brabazon and O’Neill [1] suggest considering fitness as direct profit tempered by the maximum drawdown (maximum cumulative loss experienced since the start of the trading period in question) in order to protect a trading implementation from “catastrophic loss.” In particular, return less the cost of maximum drawdown and return divided by maximum drawdown adjusted by a ratio of winning trades to overall trades are suggested in [1] (with no associated results provided for the latter). We examine fitness evaluation using the unmodified value of assets held (both USD and alternate currency), direct value minus maximum drawdown (Equation 2), and direct value simply divided by maximum drawdown (Equation 3):

$$fitness = value - maximumdrawdown, \quad (2)$$

$$fitness = value / maximumdrawdown. \quad (3)$$

Equation 2 serves to provide a fitness measure that creates more of an aversion to loss than raw profit evaluation, but is not as conservative as 3. We also examined the scaling of both the value of assets and Equation 2 by a win ratio, as originally suggested in [1]. The win ratio was the number of successful trades (gain or no loss of overall value) conducted over all trades considered by our moving window. However, due to the nature of our implementation, the win ratio simply served to reinforce the same trading decisions made by simple valuation of assets based on the cumulative trading evaluations at each segment of the moving window. In other words, both simple valuation and simple valuation adjusted by win ratio—for our particular implementation—generated near-identical results for the data series. The adjustment of equation (Equation 3) by a win ratio yielded, with a high level of consistency, less profitable results across all exchange rate series. In light of these considerations, we present results for fitness evaluation using simple valuation and Equations 2 and 3.

5. FOREX TRADING RESULTS

The worth of the assets held by the live trading system are evaluated for each of 365 days of exchange rates, with the initial 16 values withheld for seeding of technical indicators. The rates examined were CAD/USD, EUR/USD, GBP/USD, and JPY/USD from August 5, 2008 to August 5, 2009 using daily Bank of Canada [13] nominal noon exchange rates. This time period was arbitrarily chosen but includes a period of global market correction to demonstrate robustness of the algorithm. Variation across trials was practically non-existent or negligible, so a single arbitrary run is shown. LGP implementations using the fitness types described in the previous section are evaluated in terms of

Table 1: Trading Activity Summary

	CAD	EUR	GBP	JPY
<i>Trading Opportunities not Taken (%)</i>				
Raw Fitness	30.61	35.78	40.95	28.44
Moderate	28.88	32.33	35.78	28.87
Conservative	37.93	28.88	38.79	28.87
<i>Max/Min Invested (%)</i>				
Raw Fitness	30.60	35.77	39.66	21.98
Moderate	28.88	31.90	35.34	25.43
Conservative	35.34	26.29	33.19	24.57

performance trends over time, trading strategies, and final profit in the subsections that follow.

5.1 Performance Trends

Firstly, we examine the ability of the algorithm to trade over time using the value of its currencies held. Starting trading with \$1000, the total worth (total value of both non-USD currency and USD currency) of the live trading system in the base (USD) currency for the three fitness functions described in the previous section are provided in Figure 2. The value of a total initial investment of funds in the non-USD currency net interest differential is indicated as *Buy and Hold* (and is a direct reflection of the underlying rate), unmodified value of assets held is labeled *Raw Fitness*, and Equations 2 and 3 are denoted *Moderate* and *Conservative*, respectively, to reflect their approximate aversion to risk.

Is it evident from CAD/USD and GBP/USD in Figure 2 that the conservative fitness measure achieves the greatest profit when there are not extreme changes in value of the foreign (non-USD) currency. In these cases, the other fitness measures slightly lag behind or outperform buy-and-hold. When the rate volatility involves larger changes over shorter periods, as in EUR/USD (Figure 2), a moderate adjustment by maximum drawdown outperforms the other fitness measures that generally lag behind buy-and-hold throughout the time period. In the case of fairly consistent climb in value of the foreign currency in relation to USD, all fitness types work to catch up while anticipating possible indications of decline, leading to close or sub-par performance compared to buy and hold (JPY/USD, Figure 2).

5.2 Trading Analysis

All of these performance results are a direct reflection of the underlying trading activity of the implementations, which we now examine. To examine the dynamic trading strategy of the implementations, we track the percentage of total worth of each implementation that is kept in the non-USD currency at any given time over the time series. To interpret these figures keep in mind that if the algorithm anticipates an opportunity for profit, a higher percentage of USD will be placed in the non-USD currency to be sold later at a profit. These results are shown below in Figure 3. The actual percentage of trading opportunities not taken and percentage of trading opportunities where the maximum or minimum was invested in non-USD currency over the entire time period is summarized in Table 1.

It is evident from Table 1 that all fitness types caused the LGP, in general, to typically stop trading only when it was completely in or out of the market (proportion of trading

opportunities not taken equals proportion of trading where the maximum or minimum was invested). Each fitness type did have slightly different trading strategies: In the largely downward trending GBP/USD market (see Figure 2) the conservative fitness type holds a partially invested position on approximately 5.6% of days (proportion of trading opportunities not taken less proportion of trading where the maximum or minimum was invested). In the initially upward trending JPY/USD market (see Figure 2), the raw fitness implementation holds a partially invested position for 6.5% of days. To summarize, the systems are very seldom partially invested when not trading, but partial investment can occur when fitness is conservative in a downward market and when fitness is most aggressive in an upward market. Overall, Table 1 indicates moderate levels of trading where implementations do not act on approximately 29% and 39% of available trades. Figure 3 (following page) looks complex at first glance but contains easily discernible information about underlying trading strategies if the pattern of spikes indicating assets in and out of the market are examined. As expected, the number of actions taken by the conservative implementation is the lowest, followed by moderate and raw fitness in that order. This can be seen by simply examining the number of spikes that each implementation causes from a fully invested position (a position of approximately 100% non-USD investment). The conservative fitness type tends to take a trading action later than the other fitness types, but it will do so with a larger amount of capital (as seen by the general magnitude of spikes).

5.3 Profitability Analysis

The relative performance of each algorithm based on the underlying trading strategy just examined is shown succinctly in Figures 4 to 7, where the cumulative performance of all fitness types are compared to that of buy-and-hold over the entire trading period. Bottom, middle, and top of boxes indicate lower quartile, median, and upper quartile values, respectively. If notches of boxes do not overlap, medians of the two sets of data differ at the 0.95 confidence interval. Points are outliers to whiskers of 1.5 times the interquartile range. The symbol ‘+’ denotes points from 1.5 to 3 times the interquartile range, and ‘o’ represents points outside 3 times the interquartile range.

From Figures 4 to 7, it is evident that the conservative fitness outperformed (95% statistical significance) all other fitness types in three of the four currency pairs we considered. Also, when conservative fitness outperforms the other measures, it does so by a considerable margin. It appears that allowing an increased impact of maximum drawdown on the fitness function (in particular, *Conservative* over *Moderate* or *Raw Fitness*) is beneficial (with statistical significance) for currency value data that show sustained loss against the USD with some volatility (CAD/USD and GBP/USD). Figure 5 also shows that a moderate impact of maximum drawdown is more appropriate for large value swings in the exchange rate for EUR/USD. Table 2 provides the proportion of profitable buys, protective sells, and trades conducted (out of all possible trading opportunities). Many works consider proportion of profitable trades as a measure of success for an algorithm applied to market analysis. However, as noted in [1], this measure does not even provide insight into actual profit generated. One reason for this is that many trades help overall profitability by selling to prevent further




Figure 2: Mean total value of non-USD and USD currencies held given initial \$1000 USD.




Figure 3: Percentage of trading system total worth held in non-USD currency.

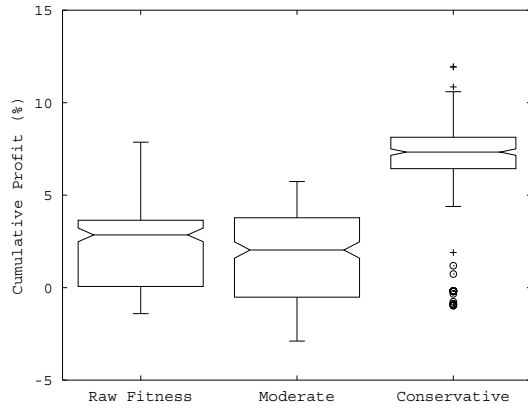


Figure 4: Cumulative profit (%) greater than Buy and Hold for CAD/USD.

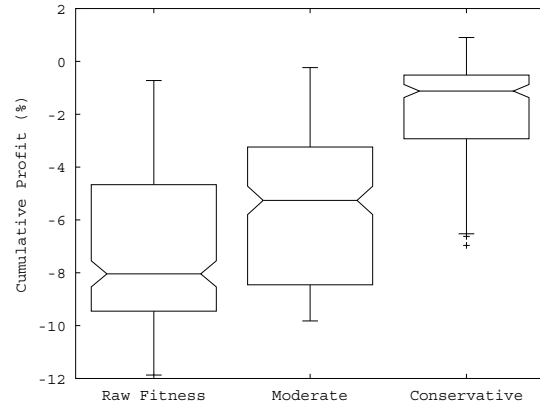


Figure 7: Cumulative profit (%) greater than Buy and Hold for JPY/USD.




Figure 5: Cumulative profit (%) greater than Buy and Hold for EUR/USD.

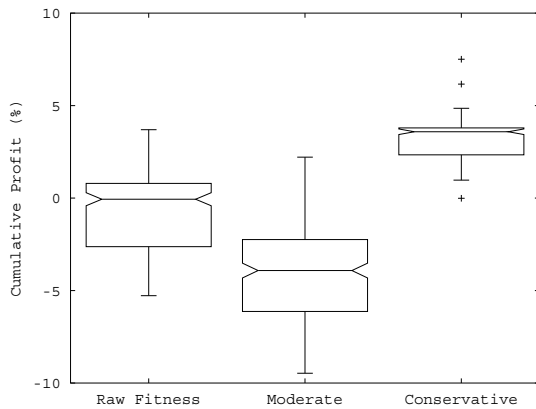


Figure 6: Cumulative profit (%) greater than Buy and Hold for GBP/USD.

Table 2: Trading Profitability Summary

	CAD	EUR	GBP	JPY
<i>Profitable Buys (%)</i>				
Raw Fitness	100.00	93.75	95.89	85.06
Moderate	98.84	96.63	87.34	92.31
Conservative	97.44	97.67	96.15	89.66
<i>Protective Sells (%)</i>				
Raw Fitness	90.14	92.19	95.24	94.87
Moderate	97.44	92.54	95.65	95.89
Conservative	92.31	93.59	93.65	90.91
<i>Final Profit (%)</i>				
Raw Fitness	0.0020	-5.72	-9.13	4.27
Moderate	-0.77	2.64	13.14	5.65
Conservative	8.98	-6.59	-2.71	13.44

losses. Thus, we use the alternative metrics of *profitable buy* and *protective sells* as introduced in [16]. A profitable buy is defined as a buy where the total value of currencies held at a time prior to the next sell exceeds the total value at the time of purchase (less slippage and interest differential). Similarly, a protective sell prevents further losses and is defined as a sell where the total value of currencies held at a time prior to the next buy is less than the total value at the time of sale (less slippage and interest differential).

Table 2 shows very impressive accuracy for both profitable buys and protective sells for all fitness types in the LGP system. Profitable buys range from approximately 87% to 100% and protective sells ranged in accuracy from 90% to 97% across all currency pairs and fitness types. However, there was no noticeable trend in terms of the accuracy of either profitables buys or protective sells across any type of fitness. That is, no particular fitness type consistently produced more accurate buying or selling compared to any other (even for the similar trend types of CAD/USD and GBP/USD where the conservative fitness measure dominated the others, see Figures 4 and 6).

In terms of profits generated over the selected time period, some currencies did not yield final profit when traded against the USD; however, positive returns ranged from profits of

0.002% to 13.44%. Comparable studies of other foreign exchange trading systems in the literature (using interday, non-leveraged rates over periods ranging from a partial year up to two years) yielded final profits of 1 to 6% [11], 0.1 to 5% [1], and 6.5% [9]. Our results are thus promising; however, it should be noted that final profit depends heavily on the trading period chosen (especially the arbitrary stopping point) and is not completely comparable across studies.

6. CONCLUSIONS

We examined the application of a linear genetic programming (LGP) system to currency trading in the foreign exchange market, examining the effect of modifications to the fitness function to provide varying degrees of investing conservatism. For the implementation discussed here, which evolves solutions for multiple prior trading days to decide on an action for the immediately proceeding day, an emphasis on ratio of winning trades had little or no impact. The greatest emphasis on maximum drawdown generated higher profits than raw fitness or moderate emphasis on maximum drawdown for most trend types. However, moderate fitness seemed to perform better on the trend with large value changes with high volatility. The overall final profits generated on the time series were competitive with other similar studies. The LGP system, regardless of fitness type, generated a very impressive proportion of profitable buys (85% to 100%) and protective sells (90 to 97%) with moderate levels of trading (29 to 39% of possible trading opportunities not taken). Future work will examine the success of LGP on other types of currency volatility patterns.

7. ACKNOWLEDGMENTS

WB acknowledges funding from NSERC under the Discovery Grant Program RGPIN 283304-07 and from Canadian Foundation for Innovation under CFI 204503.

8. REFERENCES

- [1] A. Brabazon and M. O'Neill. *Biologically Inspired Algorithms for Financial Modeling*. Springer Verlag, Berlin, 2006.
- [2] M. Brameier and W. Banzhaf. *Linear Genetic Programming*. Springer, New York, 2007.
- [3] M. Dacorogna, R. Gençay, U. Muller, R. Olsen, and O. Picket. *An Introduction to High-Frequency Finance*. Academic Press, San Diego, 2001.
- [4] M. A. H. Dempster and C. M. Jones. A real-time adaptive trading system using genetic programming. *Quantitative Finance*, 1:397–413, 2000.
- [5] M. A. H. Dempster, T. W. Payne, Y. Romahi, and G. W. P. Thompson. Computational learning techniques for intraday FX trading using popular technical indicators. *IEEE Transactions on Neural Networks*, 12(4):744–754, July 2001.
- [6] M. A. H. Dempster and Y. S. Romahi. Intraday fx trading: An evolutionary reinforcement learning approach. In *IDEAL '02: Proceedings of the Third International Conference on Intelligent Data Engineering and Automated Learning*, pages 347–358, London, UK, 2002. Springer-Verlag.
- [7] C. Faith. *Way of the Turtle: The Secret Methods that Turned Ordinary People into Legendary Traders*. McGraw-Hill, New York, 2007.
- [8] A. Hirabayashi, C. Aranha, and H. Iba. Optimization of the trading rule in foreign exchange using genetic algorithm. In *GECCO 2009 Conference Proceedings*, pages 1529–1536. ACM Press, July 2009.
- [9] A. Hryshko and T. Downs. System for foreign exchange trading using genetic algorithms and reinforcement learning. *Intern. J. Syst. Sci.*, 35(13-14):763–774, 2004.
- [10] C. J. Neely and P. A. Weller. Intraday technical trading in the foreign exchange market. *Journal of International Money and Finance*, 22(2):223 – 237, 2003.
- [11] C. J. Neely, P. A. Weller, and R. Dittmar. Is technical analysis in the foreign exchange market profitable? A genetic programming approach. *The Journal of Financial and Quantitative Analysis*, 32(4):405–426, Dec. 1997.
- [12] C. J. Neely, P. A. Weller, and R. Dittmar. The adaptive market hypothesis: Evidence from the foreign exchange market. *The Journal of Financial and Quantitative Analysis*, 44(2):223–237, April 2009.
- [13] Bank of Canada. <http://www.bank-banque-canada.ca>.
- [14] P. Saks and D. Maringer. Evolutionary money management. In *EvoWorkshops '09: Proceedings of the EvoWorkshops 2009 on Applications of Evolutionary Computing*, pages 162–171, Berlin, Heidelberg, 2009. Springer-Verlag.
- [15] R. Schwaerzel and T. Bylander. Predicting currency exchange rates by genetic programming with trigonometric functions and high-order statistics. In M. K. et al., editor, *GECCO 2006: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, volume 1, pages 955–956, Seattle, Washington, USA, 8-12 July 2006. ACM Press.
- [16] G. Wilson and W. Banzhaf. Prediction of interday stock prices using developmental and linear genetic programming. In *EvoWorkshops '09: Proceedings of the EvoWorkshops 2009 on Applications of Evolutionary Computing*, pages 172–181, Berlin, Heidelberg, 2009. Springer-Verlag.