A GENETIC PROGRAMMING APPROACH TO MODEL INTERNATIONAL SHORT-TERM CAPITAL FLOW

Tina Yu, Shu-Heng Chen and Tzu-Wen Kuo

ABSTRACT

We model international short-term capital flow by identifying technical trading rules in short-term capital markets using Genetic Programming (GP). The simulation results suggest that the international short-term markets was quite efficient during the period of 1997–2002, with most GP generated trading strategies recommending buy-and-hold on one or two assets. The out-of-sample performance of GP trading strategies varies from year to year. However, many of the strategies are able to forecast Taiwan stock market down time and avoid making futile investment. Investigation of Automatically Defined Functions shows that they do not give advantages or disadvantages to the GP results.

1. INTRODUCTION

Hot money, or speculative capital, is raising some concerns in Chinese economy. During the first half of this year (2003), about US$25 billion in short-term speculative funds sneaked into China as investors bet on possible sharp appreciation

Applications of Artificial Intelligence in Finance and Economics
Advances in Econometrics, Volume 19, 45–70

ISSN: 0731-9053/doi:10.1016/S0731-9053(04)19002-6

of the local currency *Renminbi*. Speculative capital mostly flows into areas with high liquidity, such as the security and bonds markets, as it is for short-term investments. Without being invested in industries, this money usually does not damage the overall economy once it is withdrawn.[1] Nevertheless, Chinese government has to take heed of possible longer-term fallout from speculation.

Unlike the normal direct investment, speculative capital moves very quickly among international capital markets, sometimes with very huge amount (as the Asian Crisis has demonstrated). Therefore, it can be always a potential threat for macroeconomic stability. If we can predicate the short-term capital movements, it becomes possible to control and to stabilize the economy under the influence of hot money.

In short-term international capital movements, technical trading rules play an important role as they reveal investors' behavior. This work models international short-term capital flow by identifying technical trading rules in short-term capital markets. Through the simulation, we investigate if there exists trading strategies that are capable of predicting the capital inflow and out-flow, hence make profitable investment. The modeling and simulation were conducted using Genetic Programming (GP) (Koza, 1992), a novel approach for this task. Its effectiveness will be analyzed and discussed.

As a first step, we use Taiwan as the host country and model the short-term capital flow between Taiwan and four other foreign countries: United States, Hong Kong, Japan and United Kingdom. In other words, the speculator resides in Taiwan, investing Taiwan currency to other foreign assets to pursue the highest returns.

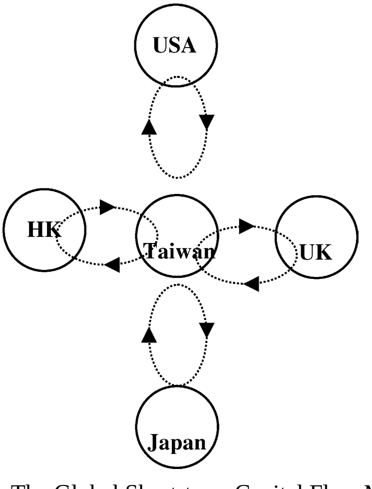

Fig. 1. The Global Short-term Capital Flow Model.

The two types of short-term assets considered here are currency and stocks, whose transactions are governed by stock markets and foreign exchange markets. This overall model gives a global picture of the short-term capital inflow and outflow between Taiwan and four foreign countries (see Fig. 1).

The paper is organized as follows. Section 2 gives the background of this work. It explains technical analysis in financial markets and surveys the applications of GP to model financial trading strategies. Section 3 describes the capital flow model representation and Section 4 gives the GP trading strategy structure. The financial data used for modeling and simulation are explained in Section 5. Section 6 gives the GP experimental setup. In Section 7, the benchmark used to evaluate GP trading strategies is explained. Section 8 presents the experimental results. The analysis of GP trading strategies is presented in Section 9. Finally, Section 10 gives the concluding remarks and outlines the direction of future work.

2. BACKGROUND

One driving force of short-term capital movement is the opportunities of profit. The prediction of short-term capital flow can therefore be viewed as the forecast of positive investment returns. One empirical approach to identify profitable capital trading is technical analysis. This approach uses historical price information to study price trends. This technique was originated from the work of Charles Dow in the late 1800 and is now widely used by investment professionals as inputs for trading decisions (Pring, 1991).

Based on technical analysis techniques, various trading rules have been developed. Examples include *moving average, filter* and *trading-range break* (see Section 4.2 for more explanation). In (Brock et al., 1992), they reported that *moving average* and *trading-range break* give significant positive returns on Dow Jones Index from 1897 to 1986. Similarly, Cooper (1999) showed that *filter* strategy can out-perform buy-and-hold under relatively low transaction cost on NYSE and AMEX stocks for the 1962–1993 period. These studies are encouraging evidences indicating that it is possible to devise profitable trading rules for financial markets.

However, one concern toward these studies is that the investigated trading rules are decided *ex post*. It is possible that the selected trading rule is favored by the tested time periods. If the investor has to make a choice about what rule or combination of rules to use at the beginning of the sample period, the reported returns may have not occurred. In order to obtain true out-of-sample performance, GP has been used to derive the trading rules for analysis (Allen & Karjalainen, 1999; Neely et al., 1997, 1999; Wang, 2000).

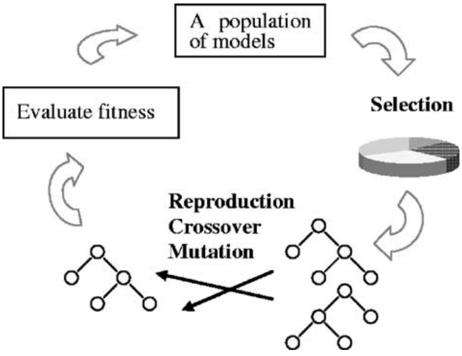

Fig. 2. Genetic Programming Cycle.

2.1. Genetic Programming

GP is a population-based search algorithm developed by John Koza (1992). It mimics the process of natural evolution to search for optimal solutions of a given problem. Figure 2 depicts the GP process cycle. Initially, a population of models is randomly created. Based on their fitness, better models are selected for reproduction. Using alteration operations, such as crossover and mutation, new offspring are generated to form a new generation. This process of selection, alteration and fitness evaluation continues until the specified termination criterion is met. The best model at the end of the process is the final model.

Various representations, selection and alteration schemes have been proposed to suit different applications. In this work, the model is represented as a parse tree that is evaluated to give trading decisions. The financial return after executing the trading decisions becomes the fitness of the model. Section 4 gives more details on the structure of GP trading strategies.

2.2. Related Works

Targeted toward different financial markets, different researchers have applied GP to generate trading rules and to analyze their profitability. For example, Allen and Karjalainen (1999) studied S&P 500 index from 1928 to 1995. They reported that the evolved GP trading rules do not earn consistent excess returns over buy-and-hold after the transaction costs. In contrast, Neely et al. (1997) reported that their GP trading rules for foreign exchange markets were able to gain excess returns for six exchange rates over the period of 1981–1995. Wang

(2000) suggested that this conflicting result might be due to the fact that foreign exchange markets have a lower transaction cost than the trading cost in the S&P index stock market. Another reason Wang suggested is that Neely et al. did not use the rolling forward approach to test their results for different time periods while Allen and Karjalainen did (see Section 5 for the explanation of rolling forward approach). Finally, Wang pointed out that these two works used different benchmarks to assess their GP trading rules: Allen and Karjalainen used the return from buy-and-hold while Neely et al. used zero return, because there is no well-defined buy-and-hold strategy in the foreign exchange markets.

Using a similar GP setup as that of Allen and Karjalainen (1999), Wang (2000) also investigated GP rules to trade in S&P 500 futures markets alone and to trade in both S&P 500 spot and futures markets simultaneously. He reported that GP trading rules are not able to beat buy-and-hold in both cases. Additionally, he also incorporated Automatically Defined Function (ADF) in his GP experiments. He reported that ADFs made the representation of the trading rules simpler by avoiding duplication of the same branches. In his work, Wang did not compare the results from GP with the results from ADF-GP.

Similar to the trading model of Wang, our short-term capital flow model allows trading in two kinds of financial markets (stock and foreign exchange) simultaneously. Moreover, we also included ADFs in our GP implementations. However, the implementations of our ADFs have more variation than that of Wang's. We also used a different data transformation method to normalize time series. Consequently, the evolved GP trading rules have different interpretations (see Section 9.1).

There are other works using Genetic Algorithms (GA) and/or Neural Network (NN) to make investment decisions. For example, Kassicieh et al. (1997) applied GA to determine the time to trade in different financial markets by selecting a subset of 10 given economic indicator time series. Baba et al. (2000) applied GA/NN hybrid to devise their decision support system for trading in Tokyo stock markets. Although GA and NN are powerful modeling tools, we find GP more suitable for our work because it has a natural representation (S-expression) for modeling trading rules. If we use GA or NN, there is an inevitable extra task of mapping the GA and NN structures to the technical trading rules.

3. MODEL REPRESENTATION

The representation of our short-term capital flow model between Taiwan and a foreign country is a directed graph. Each node in the graph represents an asset. For example, Fig. 3 gives the capital flow model between Taiwan and United States.

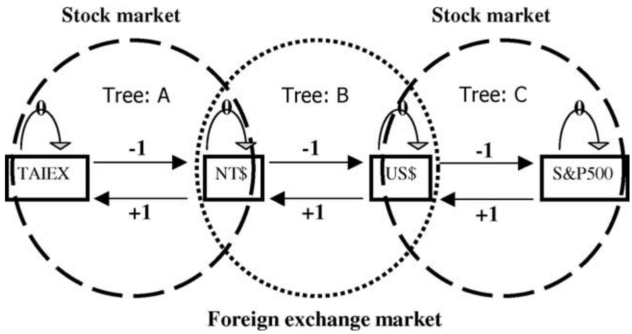

Fig. 3. The Short-term Capital Flow Model Represented as a Directed Graph.

From the left to the right, the four nodes represent Taiwan stock market (TAIEX), Taiwan currency (NT$), United States currency (US$), and United States stock market (S&P 500). This model encompasses three capital markets: Taiwan stock market, U.S. stock market and Taiwan-U.S. currency exchange market.

Funds in one asset can be transferred into one or more other assets, through the transactions in the related capital markets. For example, funds in NT$ may be changed to US$ by trading in Taiwan foreign exchange market. One can also use them to buy Taiwan stocks in the Taiwan stock market.

Initially, the fund is placed in foreign currency. At each time step, the fund may be reallocated to other assets, according to the trading decisions made for the three capital markets. These three decisions made up the trading strategies to be carried out by an investor. More details on GP trading strategies are given in the following section.

A trading decision may be to buy an asset, to sell an asset or to do nothing. For the purpose of generality, we structure a financial market with two assets, one on left and one on right. When the decision is to transfer a fund from the asset on the right to the asset on the left, a "$+1$" is signaled. When the decision is to transfer a fund from the asset on the left to the one on the right, a "-1" is signaled. Signal "0" means do nothing. Table 1 gives the 27 possible combinations of trading decisions. Assuming at time t, the fund in TAIEX is A, in NT$ is B, in US$ is C and in S&P500 is D, the table gives the fund allocations at time $t+1$.

When the decision is to trade ($+1$ or -1), half of the current fund is transferred to the designated asset. For example, if the trading strategy is $\{-1, -1, -1\}$, half of the TAIEX funds (A) will be moved to NT$; half of the original NT$ fund (B) will be moved to US$ and half of the original US$ fund (C) will be moved to S&P500. A trading strategy may cause an original fund to be completely transferred out, e.g. $\{1, -1, -1\}$ trades all NT$ with TAIEX and US$. However, the maximum amount

Table 1. Trading Decisions and Their Funds Reallocation Results.

TSM	CEM	FSM	$TAIEX_{t+1}$	$NT\$_{t+1}$	$U.S.\$_{t+1}$	$S\&P500_{t+1}$
−1	−1	−1	0.5A	0.5A + 0.5B	0.5B + 0.5C	0.5C + D
−1	−1	0	0.5A	0.5A + 0.5B	0.5B + C	D
−1	−1	1	0.5A	0.5A + 0.5B	0.5B + C + 0.5D	0.5D
−1	0	−1	0.5A	0.5A + B	0.5C	0.5C + D
−1	0	0	0.5A	0.5A + B	C	D
−1	0	1	0.5A	0.5A + B	C + 0.5D	0.5D
−1	1	−1	0.5A	0.5A + B + 0.5C	0	0.5C + D
−1	1	0	0.5A	0.5A + B + 0.5C	0.5C	D
−1	1	1	0.5A	0.5A + B + 0.5C	0.5C + 0.5D	0.5D
0	−1	−1	A	0.5B	0.5B + 0.5C	0.5C + D
0	−1	0	A	0.5B	0.5B + C	D
0	−1	1	A	0.5B	0.5B + C + 0.5D	0.5D
0	0	−1	A	B	0.5C	0.5C + D
0	0	0	A	B	C	D
0	0	1	A	B	C + 0.5D	0.5D
0	1	−1	A	B + 0.5C	0	0.5C + D
0	1	0	A	B + 0.5C	0.5C	D
0	1	1	A	B + 0.5C	0.5C + 0.5D	0.5D
1	−1	−1	A + 0.5B	0	0.5B + 0.5C	0.5C + D
1	−1	0	A + 0.5B	0	0.5B + C	D
1	−1	1	A + 0.5B	0	0.5B + C + 0.5D	0.5D
1	0	−1	A + 0.5B	0.5B	0.5C	0.5C + D
1	0	0	A + 0.5B	0.5B	C	D
1	0	1	A + 0.5B	0.5B	C + 0.5D	0.5D
1	1	−1	A + 0.5B	0.5B + 0.5C	0	0.5C + D
1	1	0	A + 0.5B	0.5B + 0.5C	0.5C	D
1	1	1	A + 0.5B	0.5B + 0.5C	0.5C + 0.5D	0.5D

Note: TSM: Taiwan Stock Market; CEM: Currency Exchange Market; FSM: Foreign Stock Market. The table is simplified in that no transaction cost is considered. The modeling process, however, does take transaction cost into account.

of fund that one asset can acquire is half of its two neighboring assets. For example, the trading strategy $\{0, -1, 1\}$ leads to an increase of US$ by half of the original NT$ fund and half of the original S&P500 fund. This is a rather conservative setup. We therefore adjust the modeling and simulation procedure to execute a transaction multiple (10) times in one time step,[2] with the transaction cost charged once only. This leads to almost 100% of the original fund in one asset to be transferred to the designated asset in one time step. Nevertheless, to reallocate all the funds $(A + B + C + D)$ into single asset, it still requires at least three time steps.

To be close to the reality, the model does not allow direct capital flow between international stocks. In the real world, the trading between stocks in two different

Table 2. Trading Decision Table.

Rule 1 Recommendation	Rule 2 Recommendation	Final Decision
True	False	+1
False	True	−1
True	True	0
False	False	0

countries requires an intermediate step of currency exchange. For example, to trade a Taiwan stock with a U.S. stock, the Taiwan stock has to be cashed into Taiwan currency, which is exchanged to U.S. currency, which is then used to purchase the U.S. stock.

4. GP TRADING STRATEGIES

A GP trading strategy consists of three trading decisions made for the three financial markets. Each trading decision (+1, −1 or 0) is determined by a pair of GP rules. The first rule decides whether to move funds from the right asset to the left asset (True) or not (False). The second rule decides whether to move funds from the left asset to the right asset (True) or not (False). The final decision is derived according to Table 2.

A GP rule has a tree structure. Figure 4 gives a trading rule example. It says, "If the 15-day moving average is greater than the 250-day moving average, then trade. Otherwise, if the closing exchange rate has risen by more than 1% above its minimum over the previous 10 days, then trade. Otherwise, do not trade."

With three trading decisions, each is determined by two rules; a GP trading strategy consists of six GP trees. Figure 5 gives the structure of a GP trading strategy. Note that the labels Tree-A, Tree-B and Tree-C correspond to those in Fig. 3.

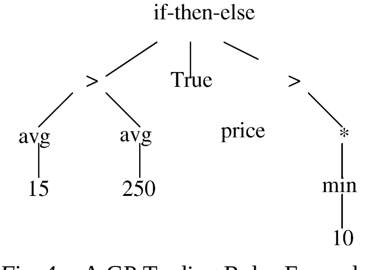

Fig. 4. A GP Trading Rules Example.

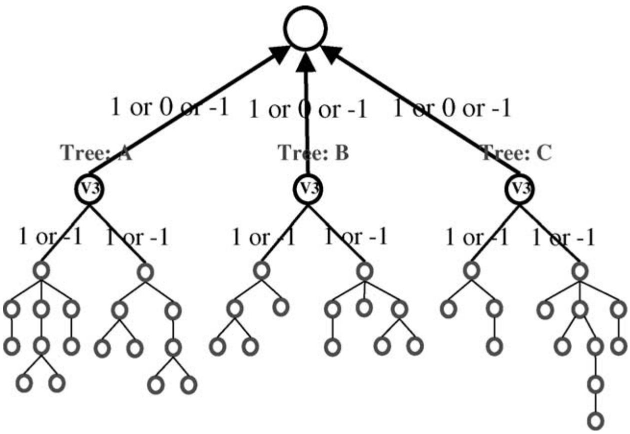

Fig. 5. The GP Trading Strategy Structure.

The following functions are provided to construct the internal nodes of a GP tree:

- Boolean function: *and, or, not,* $<$, $>$, *if-then-else*
- Numerical function: $+$, $-$, $\times$, $\div$, *average, max, min, norm, lag*

The function *average* computes the moving average of a variable in a time window specified by the integer argument. For example, *average* (x, 250) at time t is the arithmetic mean of $x_{t-1}, x_{t-2}, \ldots, x_{t-250}$. The function *max* returns the largest value of a variable during a time window specified by the integer argument. For example, *max* (y, 3) at time t is equivalent to $max(y_{t-1}, y_{t-2}, y_{t-3})$. Similarly, the function *min* returns the smallest value of a variable during a time window specified by the integer argument. The function *norm* computes the absolute value of the given real number. The function *lag* returns the value of a variable lagged by a number of days specified by the integer argument. For example, *lag* (z, 3) at time t is z_{t-3}. These functions are commonly used by financial traders to decide their trading strategies, hence are reasonable building blocks for GP to construct trading rules.

GP tree leaf nodes can be a value from the following three types of terminals:

- Input variables: $\mathrm{TW_{IR}}$, $\mathrm{TW_{SI}}$, $\mathrm{FC_{IR}}$, $\mathrm{FC_{SI}}$, NTD/FD
- Numerical constants: 100 constants randomly generated between 0.0 and 10.0
- Boolean constants: True, False

Input variables include: *interest rate* in Taiwan ($\mathrm{TW_{IR}}$) and the foreign country ($\mathrm{FC_{IR}}$); *stock index* in Taiwan stock market ($\mathrm{TW_{SI}}$) and the foreign country stock market ($\mathrm{FC_{SI}}$); the *exchange rate* between Taiwan and the foreign country (NTD/FD). These financial time series will be explained with more details in

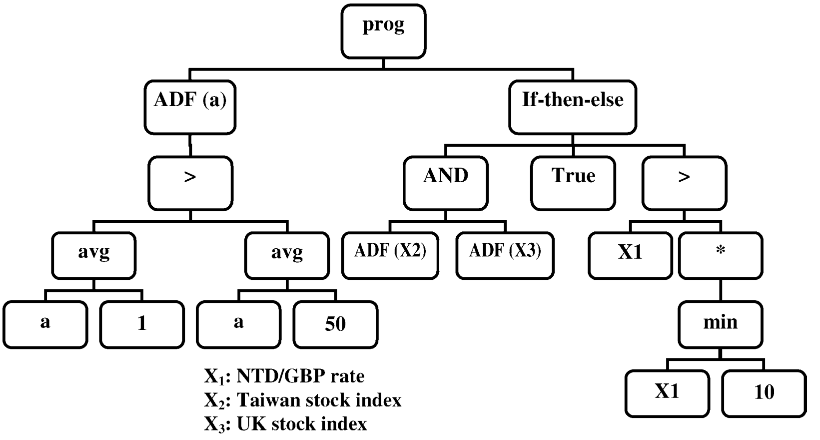

Fig. 6. An ADF-GP Trading Strategy Example.

Section 5. Real-valued constants may be truncated into integer value if they are passed over to time series functions, such as *lag*.

4.1. Automatically Defined Functions

Automatically Defined Function (ADF) is a mechanism devised by Koza to extend GP ability to solve problems with regularity, symmetry and homogeneity (Koza, 1994). ADFs are subroutines that are simultaneously evolved with the GP main programs. Figure 6 gives an example GP trading rule with one ADF. The left branch of the tree is an ADF while the right branch is the main trading rule. The ADF takes one argument (a time series variable) and checks if its 1-day moving average is greater than its 50-day moving average. This ADF is called twice in the GP main trading rule: ADF($\times$2) takes Taiwan stock index as the argument while ADF($\times$3) takes UK stock index as the argument.

An ADF is evolved simultaneously with the GP main trading rule. If a trading rule contains patterns, ADF-GP may discover and extract them as ADFs, which are then called from the GP main trading rule. We implemented ADFs in three different ways for three different purposes:

- One ADF is included in each trading strategy. This investigates whether regularity exists in profitable trading strategies and as to whether GP is able to discover them. Since the time series are transformed by dividing them by 250-day moving average (see Section 5), ADF is used to identify patterns in the *change of trend* that provide profitable trading.

- One partially defined ADF is included in each trading strategy. The ADF is initially seeded with one of the commonly used technical trading rules (see Section 4.2). They are then evolved during the GP runs. With the transformed time series, this implementation is to discover if the provided technical trading rules (and their variations) are effective on the transformed time series data.
- Three partially defined ADFs are included in each trading strategy. This is the same as the above except three, instead of one, ADFs are used.

The function and terminal sets used to evolve ADFs are the same as that used to evolve the GP main program. For ADF-GP, an extra function (the name of the ADF) is included in the GP main program function set.

4.2. Technical Trading Rules

Two types of technical trading rules are provided for GP to initialize its ADFs: *moving average rules* and *filter rules*. Moving average rules include a class of rules where the trading signals are decided by comparing a short-run with a long-run moving average in the same time series, producing a "buy" signal when the short-run moving average cuts the long-run moving average from below. This rule can be implemented in many different ways by specifying different short and long periods. We have included the following five implementations: (1–50), (1–150), (5–150), (1–200), and (2–200), where the first number is the short while the second number indicates the long. We also implemented a band moving average rule, where the band is 0.01, i.e. signal "buy" if the short-run moving average exceeds the long-run moving average by 1%.

Filter rules include a class of trading rules where the trading signals are decided by comparing the current price with its local low or with its local high over a past period of time. We select three time lengths (50, 150, 200) to implement this class of rules. We also implemented two band filter rules, one with band 0.01 and the other with band -0.01. In the first case, a "buy" signal is generated if the current price exceeds the local high by 1%. In the second scenario, a "sell" signal is generated if the current price is below the local low by 1%.

Since these predefined ADFs are evolved, the final versions have different semantics and are not to be called the same names anymore.

5. DATA SET

We have acquired financial time series data for five countries (Taiwan, United States, United Kingdom, Japan and Hong Kong) between January 1, 1992 and

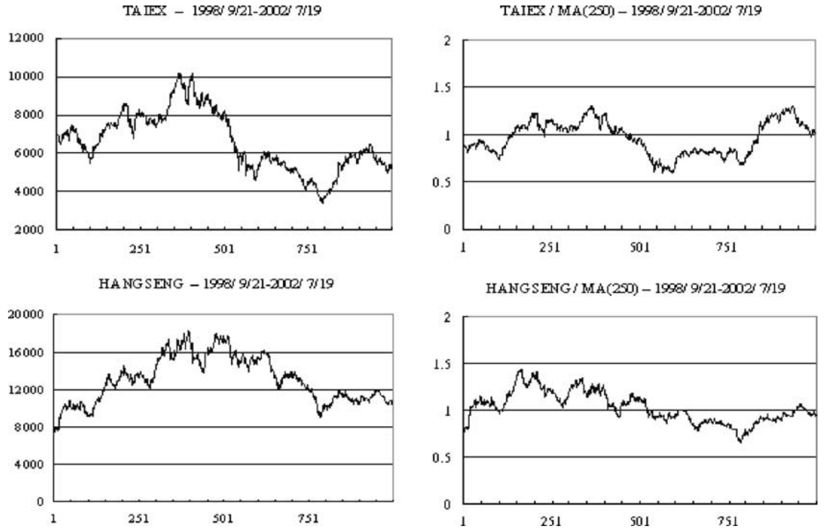

Fig. 7. Time Series Data Before and After Normalization.

December 31, 2002 from *Datastream*. The time series include: TW_{IR}, TW_{SI}, US_{IR}, US_{SI}, UK_{IR}, UK_{SI}, HK_{IR}, HK_{SI}, JP_{IR}, JP_{SI}, NTD/USD, NTD/GBP, NTD/JPY, NTD/HKD. Five times series are used to build one model. For example, TW_{IR}, TW_{SI}, US_{IR}, US_{SI} and NTD/USD are used to model Taiwan-U.S. capital flow.

Since the original time series are non-stationary, we transform them by dividing the daily data by a 250-day moving average. This is the method used by Allen and Karjalainen (1999) and Neely et al. (1997, 1999). The adjusted data oscillate around 1 and make the modeling task easier. Figure 7 gives two examples. On the left side are the two original series while on the right are the transformed ones. While the transformed series are used for modeling, the computation of GP trading strategies returns is based on the original time series. One implication of this data transformation is that GP is searching for patterns exhibited in the *change of trends* that give profitable trading strategies.

Over-fitting is an issue faced by all data modeling techniques. GP is no exception. When constructing/optimizing the trading strategies, GP tends to make the strategies producing maximum returns for the training period, which may contain noise that do not represent the overall series pattern. In order to construct trading strategies that generalize beyond the training data, we adopt two methods to run the GP experiments. The first one is to enforce parsimony pressures on the trading strategies structures, which will be discussed in Section 6. The second one is splitting the series into training, validation and out-of-sample periods. This is

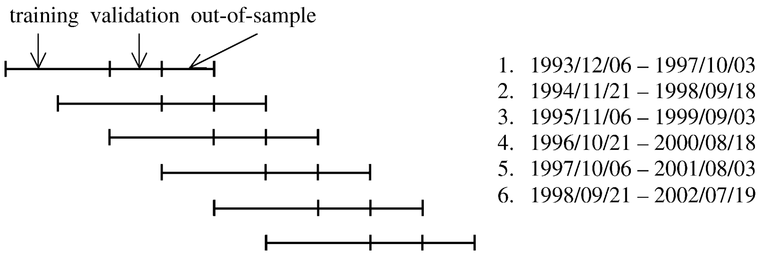

Fig. 8. Six Sequences of Time Series Data.

a commonly used approach in machine learning and data mining. We adopt the rolling forward approach first proposed by Pesaran and Timmermann (1995) and also used by Allen and Karjalainen (1999) and Wang (2000).

To start, the first 500 data (250 used to transform raw data and 250 reserved to be referred by time series function such as *lag*) were removed. This leaves 2500 data in each time series. To guard against potential data snooping in the choice of time periods, the series are organized into 6 sequences, each with 1000 data points. Among them, 500 are for training, 250 are for validation and 250 are for out-of-sample testing. The data in one series may overlap with that in other series. As shown in Fig. 8, the second half of the training period and the entire validation period at the first series are the training period at the second series. The out-of-sample testing period at the first series is the validation period at the second series. With this setup, each out-of-sample testing period is one-year (short-term) and covers a different time period.

For each data series, 20 GP runs were made. The three data periods are used in the following manner:

(1) The best trading strategy against the training period at the initial population is selected and evaluated against the validation period. This is the initial "best strategy."
(2) A new generation of trading strategies is created by recombining/modifying parts of relatively fit strategies in the previous generation.
(3) The best trading strategies against the training period at the current population is selected and evaluated against the validation period.
(4) If this strategy has a better validation fitness than the previous "best strategy", then this is considered to be the new "best strategy".
(5) Go to step 2 until the maximum number of generation is reached or there is no fitter strategy is found after a certain number of generations (a controllable parameter).
(6) The last "best strategy" is tested against the out-of-sample period. This is what we use to evaluate the performance of GP trading strategies.

In summary, the training period is used to construct/optimize GP trading strategies while the validation period is used to select the GP trading strategies, which are then applied on the out-of-sample period to give the performance of the strategies. The analysis and evaluation are based on results from the out-of-sample period.

6. EXPERIMENTAL SETUP

The control parameters used to run GP experiments are given in Table 3. We experimented with different population size (200, 500 and 1000) to run for different number of generations (100 and 200). This setup is motivated by an observation reported by Chen and Kuo (2002) that population size and the number of generations have impact on GP search efficiency when modeling chaotic time series. These results will be compared in Section 8.

The GP system is generation-based, i.e. parents do not compete with offspring for selection and reproduction. This is a less aggressive search method compared to the steady-state-based GP where the offspring are used to replace less fit individuals in the population (Syswerda, 1991). Although steady-state-based GP has the advantage that fit offspring become available for reproduction right away, there are possibilities that the population becomes converged too fast hence leads to sub-optimal solutions.

We used tournament of size 2 to select winners. This means that two individuals are randomly selected and the one with a better fitness is the winner. For crossover operation, two winners are selected. For mutation or copy operation, only one winner is needed. The new population is generated with 70% of the individuals from crossover, 10% from point mutation, 10% from tree mutation and 10% from copy operation. The best individual in the current population is always copied over to the new generation.

Table 3. Control Parameters for GP Experiments.

Parameter	Value
Population size	200, 500, 1000
Maximum generation	100, 200
Crossover rate	70%
Point mutation rate	10%
Tree mutation rate	10%
Reproduction (copy)	10%
Elite	1
Maximum tree node	50
Maximum tree depth	17

The maximum tree depth of 17 is a hard constraint that cannot be violated. A GP strategy with tree depth larger than 17 is discarded. This is necessary to accommodate the computer resources. In contrast, the maximum number of tree node (50) is a soft constraint, which is handled using penalty explained in the following section (see Yu & Bentley, 1998 for more constraint handling methods).

As mentioned in Section 5, the best rule for a training period in each generation is evaluated against validation period. If the rule has a validation fitness that is better than the previous best rule has, it is saved as the new best rule. A GP run stops if no new best rule appears for 1/4 of the specified maximum number of generations or when the maximum number of generations is reached.

The fitness of an evolved GP trading strategy is the *gross return* (R) of the investment it generates. Initially, an investment of 1 unit is made in foreign currency. At the end of the time period, its final value is the *gross return*.

To determine the fitness of a GP trading strategy, it is applied on the normalized time series to produce a series of trading decisions for the three financial markets. This decision series are executed 10 times in each time step until the end of the time period. Every time a trading decision is executed, the amount of funds in each of the four assets may change (see Table 1). Let the amount of fund transferred from A to B be $\vec{A}$, from B to C be $\vec{B}$, from C to D be $\vec{C}$, from D to C be $\overleftarrow{D}$, from C to B be $\overleftarrow{C}$, from B to A be $\overleftarrow{B}$. Also, the associated one-way transaction costs are $\mathrm{Cost}_{\mathrm{AB}}$, $\mathrm{Cost}_{\mathrm{BC}}$, $\mathrm{Cost}_{\mathrm{CD}}$, $\mathrm{Cost}_{\mathrm{DC}}$, $\mathrm{Cost}_{\mathrm{CB}}$ and $\mathrm{Cost}_{\mathrm{BA}}$. $\mathrm{TW}_{\mathrm{SI}}$ is the Taiwan stock index and $\mathrm{FC}_{\mathrm{SI}}$ is the foreign stock index. $\mathrm{TW}_{\mathrm{IR}}$ is the Taiwan currency interest rate and $\mathrm{FC}_{\mathrm{IR}}$ is the foreign currency interest rate. E is the exchange rate between the two currencies. At time $t+1$, the funds in each asset is given by:

$$
\begin{aligned}
A_{t+1} &= A_t - \vec{A} + \frac{\overleftarrow{B}}{\mathrm{TW}_{\mathrm{SI}(t)} \times (1 + \mathrm{Cost}_{\mathrm{BA}})} \\
B_{t+1} &= (B_t - \vec{B} - \overleftarrow{B}) \times (1 + \mathrm{TW}_{\mathrm{IR}(t)}) + \vec{A} \times \mathrm{TW}_{\mathrm{SI}(t)} \\
&\quad \times (1 - \mathrm{Cost}_{\mathrm{AB}}) + \overleftarrow{C} \times E_t \times (1 - \mathrm{Cost}_{\mathrm{CB}}) \\
C_{t+1} &= (C_t - \vec{C} - \overleftarrow{C}) \times (1 + \mathrm{FC}_{\mathrm{IR}(t)}) + \frac{\vec{B}}{E_t \times (1 + \mathrm{Cost}_{\mathrm{BC}})} + \overleftarrow{D} \\
&\quad \times \mathrm{FC}_{\mathrm{SI(t)}} \times (1 - \mathrm{Cost}_{\mathrm{DC}}) \\
D_{t+1} &= D_t - \overleftarrow{D} + \frac{\vec{C}}{\mathrm{FC}_{\mathrm{SI}(t)} \times (1 + \mathrm{Cost}_{\mathrm{CD}})}
\end{aligned}
$$

Different financial markets have different transaction costs. Moreover, within the same financial market, a transaction from asset A to asset B may have a different cost from that of a transaction from asset B to asset A. Table 4 gives the transaction cost implemented in this work. The costs associated with Taiwan stock market and

Table 4. Transaction Cost.

Transaction Type	Rate (%)
Cost_{AB}	0.4425
Cost_{BC}	0.2**
Cost_{CD}	0.1*
Cost_{DC}	0.43*
Cost_{CB}	0.2**
Cost_{BA}	0.1425

*Allen et al. (1999) used 0.1, 0.25 and 0.5% as the one-way transaction cost for S&P500 index market, while Wang (2000) used 0.12% for the same market.
**Neely et al. (1997) used 0.05% as the one-way transaction cost for foreign exchange markets.

Taiwan foreign exchange market are actual values. The costs associated with foreign country stock markets are estimated based on the fixed transaction tax charged to international investment and an estimated handling charge of 0.1%. Compared to the transaction cost for S&P500 stock market used by Allen and Karjalainen (1999) (0.1, 0.25 & 0.5%) and by Wang (2000) (0.12%), we have a higher transaction cost. Also, we have a higher transaction cost for foreign exchange market than that used by Neely et al. (1997) (0.05%). Normally, higher transaction costs discourage trades and reduces the number of transactions. This work intents to reflect the actual market operations, hence adapts the actual financial costs in the markets for modeling, in spite of the fact that they are higher than those used in other studies.

At the end of the time period (T), all assets are converted into the foreign currency:

$$B_{T+1} = B_T + A_T \times \text{TW}_{\text{SI}(T)} \times (1 - \text{Cost}_{\text{AB}})$$

$$C_{T+1} = C_T + \frac{B_{T+1}}{E_T \times (1 + \text{Cost}_{\text{BC}})} + D_T \times \text{FC}_{\text{SI}(T)} \times (1 - \text{Cost}_{\text{DC}})$$

The gross return is:

$$R = C_{T+1}$$

There is a penalty toward GP strategies that exceed the maximum number of 50 nodes. This soft constraint approach allows fitter strategies with a larger number of nodes to survive. Yet, it discourages tree size growth to avoid over-fitting, since trees with a large number of nodes tend to fit the training data so well that they lose their generality. The final fitness of a GP trading strategy is given in the following equation (Seshadri, 2003):

$$F = R \frac{50}{\max(\text{tree\_size},\ 50)}$$

7. BENCHMARK

The buy-and-hold (B&H) strategy is the most commonly used benchmark to evaluate financial trading strategies. With B&H, an investment made on one asset

Table 5. Return for the Buy-and-Hold Strategy.

Year	TW-U.S. Model	TW-HK Model	TW-JP Model	TW-UK Model
1997				
B&H(A)	1.2618	1.2627	1.3645	1.2431
B&H(B)	1.0187	1.0194	1.1016	1.0035
B&H(C)	1.0492	1.0536	1.0047	1.0594
B&H(D)	1.3523	1.2029	0.8240	1.3019
$R_{\mathrm{B\&H}}$	**1.1705**	**1.1346**	**1.0737**	**1.1520**
1998				
B&H(A)	0.6792	0.6805	0.7364	0.6513
B&H(B)	0.8805	0.8822	0.9546	0.8444
B&H(C)	1.0492	1.0626	1.0042	1.0688
B&H(D)	1.0432	0.5012	0.7803	0.9488
$R_{\mathrm{B\&H}}$	**0.9130**	**0.7816**	**0.8689**	**0.8783**
1999				
B&H(A)	1.2342	1.2362	1.0198	1.2862
B&H(B)	1.1340	1.1358	0.9370	1.1818
B&H(C)	1.0449	1.0492	1.0012	1.0558
B&H(D)	1.3186	1.8282	1.2897	1.2622
$R_{\mathrm{B\&H}}$	**1.1829**	**1.3124**	**1.0619**	**1.1965**
2000				
B&H(A)	1.0192	1.0249	1.0139	1.0948
B&H(B)	1.0669	1.0728	1.0612	1.1460
B&H(C)	1.0527	1.0514	1.0003	1.0541
B&H(D)	1.0933	1.2960	0.9120	1.0209
$R_{\mathrm{B\&H}}$	**1.0580**	**1.1113**	**0.9968**	**1.0790**
2001				
B&H(A)	0.4889	0.4893	0.5593	0.5101
B&H(B)	0.9351	0.9358	1.0698	0.9757
B&H(C)	1.0485	1.0494	1.0015	1.0555
B&H(D)	0.8056	0.6973	0.7592	0.8435
$R_{\mathrm{B\&H}}$	**0.8195**	**0.7930**	**0.8474**	**0.8462**
2002				
B&H(A)	1.1983	1.1991	1.1205	1.0807
B&H(B)	1.0691	1.0698	0.9997	0.9642
B&H(C)	1.0158	1.0200	1.0000	1.0406
B&H(D)	0.7024	0.8454	0.8288	0.7377
$R_{\mathrm{B\&H}}$	**0.9964**	**1.0336**	**0.9873**	**0.9558**

stays there until the end of time period. Since there are four assets in a model, the B&H strategy can be applied in four different ways: buy TAIEX and hold, buy NT$ and hold, . . ., etc. We therefore apply B&H over these four different assets. The average of their returns is used as the benchmark. Table 5 gives the B&H returns for the four different models.

8. RESULTS

For each of the four foreign countries modeled, we obtain 36 GP trading returns. These GP strategies are evolved based on 6 different data sequences using 3 different population sizes to run for 2 different numbers of generations. Each of the 36 results is the average of 20 trials. Table 6 gives the percentage of the GP trading strategies that out-performs the B&H strategy.

In TW-US, TW-JP and TW-UK models, most GP trading strategies out-perform B&H. In contrast, TW-HK model has a less number of GP trading strategies that give better returns than B&H. The number of statistically significant GP returns is given inside the parenthesis.

Different population sizes and number of generations make little difference on the GP results. For the small number of cases where they produce different results, there is not a consistent pattern showing larger (smaller) population size and/or longer (shorter) runs give better results. We checked the log files and found that most of the runs stop before generation 50 when no improved strategy on validation period was found.

Moreover, ADFs in various form, provide no improvement in performance than the standard or "vanilla" GP model we used. For those runs where vanilla GP produces better returns than B&H, ADF-GP also gives better returns. Similarly, those runs where vanilla GP produces worse returns than the B&H method, the ADF-GP performs even worse. We will analyze the ADF-GP trading strategies and give explanation of this outcome in Section 9.1.

In this section, we analyze GP trading strategies based on the vanilla GP out-of-sample results, which are summarized in Table 7. In the table, six sets of data are

Table 6. Percentage of GP Trading Strategies Results that Out-performs B&H.

GP Implementation	TW-U.S. Model	TW-HK Model	TW-JP Model	TW-UK Model
Vanilla GP	29(19)/36	8(5)/36	19(11)/36	28(21)/36
GP with 1 ADF	27(20)/36	9(3)/36	22(9)/36	26(22)/36
GP with 1 partially defined ADF	30(21)/36	13(4)/36	23(11)/36	26(20)/36
GP with 3 partially defined ADFs	29(15)/36	11(7)/36	25(8)/36	26(12)/36

Table 7. Summary of Vanilla GP Trading Strategies Results.

Year	TW-U.S. Model			TW-HK Model		
	μ	σ	t	μ	σ	t
1997	**1.2295**	0.1292	**2.0409**	1.0815	0.0753	−3.1553
	1.1764	0.0998	0.2638	1.0865	0.1015	−2.1181
	1.1397	0.1171	−1.1759	1.0784	0.0652	−3.8543
	1.1680	0.1305	−0.0846	1.1228	0.1136	−0.4636
	1.2070	0.1086	1.5040	1.1077	0.1016	−1.1834
	1.2001	0.0907	1.4604	1.0784	0.0754	−3.3333
1998	**1.0489**	0.0190	**32.0570**	0.5356	0.0755	−14.5672
	1.0287	0.0355	**14.5564**	0.5102	0.0401	−30.2427
	1.0293	0.0488	**10.6501**	0.5058	0.0200	−61.5146
	1.0437	0.0014	**428.904**	0.5035	0.0102	−121.807
	1.0243	0.0553	**9.0039**	0.5927	0.1645	−5.1336
	1.0457	0.0328	**18.0693**	0.5429	0.1481	−7.2061
1999	**1.2033**	0.0708	1.2920	1.1947	0.0765	−6.8797
	1.2344	0.0877	**2.6265**	1.2110	0.1184	−3.8310
	1.2183	0.0875	**1.8076**	1.2720	0.1424	−1.2689
	1.2146	0.0936	1.5137	1.2728	0.1225	−1.4475
	1.1845	0.1195	0.0600	1.2494	0.1283	−2.1974
	1.2585	0.1160	**2.9160**	1.2419	0.1315	−2.3980
2000	**1.0980**	0.0178	**10.0229**	1.0739	0.0688	−2.4330
	1.0915	0.0162	**9.2525**	1.0808	0.0806	−1.6928
	1.0914	0.0132	**11.3239**	1.0967	0.0842	−0.7770
	1.0779	0.0660	1.3476	1.0540	0.0508	−5.0403
	1.0821	0.0475	**2.2684**	**1.1580**	0.1255	1.6630
	1.0974	0.0336	**5.2471**	**1.1167**	0.1222	0.1965
2001	**0.8227**	0.0313	0.4588	**0.9045**	0.1477	**3.3785**
	0.8538	0.0622	**2.4679**	**0.9207**	0.1558	**3.6642**
	0.8374	0.0284	**2.8263**	**0.8779**	0.1472	**2.5781**
	0.8430	0.0655	1.6014	**0.8370**	0.1385	1.4203
	0.8425	0.0267	**3.8532**	**0.8774**	0.1741	**2.1677**
	0.8509	0.0350	**4.0088**	**0.8755**	0.1569	**2.3513**
2002	0.9516	0.1376	−1.4553	0.8649	0.0440	−17.1482
	0.9677	0.1412	−0.9081	0.8639	0.0326	−23.2644
	0.9008	0.1383	−3.0940	0.8688	0.0441	−16.7151
	0.9416	0.1589	−1.5428	0.9015	0.1078	−5.4767
	0.9203	0.1665	−2.0433	0.8860	0.0662	−9.9759
	1.0118	0.1295	0.5315	0.9195	0.0908	−5.6226

given for each of the 6 out-of-sample periods (1997–2002). Each set contains data obtained from vanilla GP runs using different combinations of population size and number of generations. The average return of 20 trials is μ; the standard deviation is σ; the *t*-statistics is *t*. Those μ values in bold are average returns which are better than the returns of B&H (see Table 5 for B&H returns). Those *t* values in bold indicate the difference between μ and the B&H return is significant at the 5% level.

As shown, the performance of GP strategies varies in different out-of-sample periods. For example, in sequence 5 period, all GP strategies out-perform B&H while in sequence 6 period, B&H gives higher returns in most of the cases. We examined time series in sequence 6 and found that both Taiwan stock and the foreign stock indices (the two most influential trading decision factors) fluctuate widely. For example, during the training period (1999 and 2000), both Taiwan stock and Hang Seng indices declined. During the validation period (2001), the markets gradually improved. However, the markets rallied during the out-of-sample testing period (2002) (see Fig. 7). As a result, the strategies trained using 1999 and 2000 periods and selected based on 2001 period are not able to perform well on 2002 period. This is a shortcoming of all machine learning techniques, including GP.

In contrast, the stock indices for training, validation and out-of-sample periods in sequence 5 have a similar pattern: the stock markets generally went down. Consequently, the strategies evolved on training period were able to perform well on the out-of-sample period. Another interesting observation is that although all markets decline in this period (with Taiwan stock market having the worst decline of 50%) and cause B&H to have low returns (see Table 5), GP strategies were able to make profitable trading decisions. Figure 9 gives two such examples. In Fig. 9(a), the 1 US$ was kept until day 151 and then invested in S&P500 when the index started rising. As a result, it has a return of 1.1131, which is better than holding it until the end of the time period (1.0485). Figure 9(b) gives a different GP strategy, which invested in S&P index stock too early and cause a negative return at the beginning. However, as the index started improving on day 151, the return

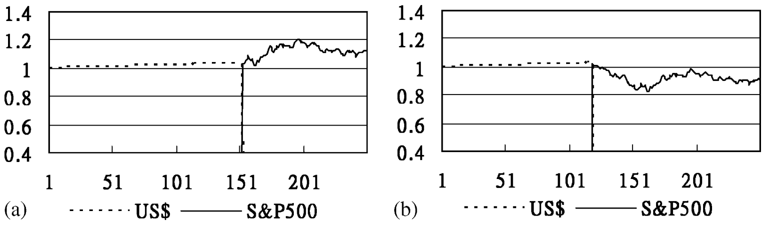

Fig. 9. Capital Flows of Two GP Trading Strategies Applied on Out-of-sample Period of 2001.

became positive. At the end of the period, the return is 0.9036, which is better than the B&H return. One important observation is that none of the GP trading strategies trained in this time sequence entered into Taiwan stock market, the worst asset to invest. This suggests GP strategies have some forecasting abilities in the sense of avoiding money-losing assets all the way to the end of the period.

The transaction frequencies in out-of-sample testing periods are mostly low: no more than 3 times in the whole year. The majority of GP strategies recommend to buy-and-hold on one or two assets. For example, for out-of-sample period 1999, GP trading strategies in TW-JP model either invest in Taiwan stock market or in Japan stock market. The first decision gives a higher return than the second decision does. There are also many strategies give zero transaction: hold the foreign currency all the way to the end of the period. Consequently, most of the GP trading strategies give returns that are close to the returns of B&H (see Tables 7 and 5). This indicates that international short-term financial markets are reasonably efficient during the years between 1997 and 2002.

Overall, the out-of-sample performance of GP trading strategies are not consistently better than that of B&H, an outcome that is consistent with the finding of Allen and Karjalainen (1999) and Wang (2000).

9. ANALYSIS OF GP TRADING STRATEGIES

9.1. Vanilla-GP Trading Strategies

Using both hard and soft constraints to enforce parsimony, the evolved GP trading strategies are not as complex as what we have expected. As mentioned in the previous section, many of them are evaluated into a simple B&H on one or two assets. These strategies either have other options blocked by constant "do nothing" decisions or recommending trading using assets which have no available fund. Overall, the decisions of GP strategies are not difficult to derive, although their financial meaning is challenging to interpret.

We have found one particular GP strategy derived from TW-UK model that provides financially meaningful recommendations. Figure 10 gives the GP tree. Without being very rigorous, the left tree can be interpreted as:

> When UK stock market bulls (suggested by a higher FTSE100 index 250-moving average than US/GBP 250-moving average exchange rate), sell Taiwan stocks to obtain Taiwan currency.

The middle tree can be interpreted as:

> Trade British Ponds with Taiwan Dollars when the exchange rate is less than the 250-day moving average.

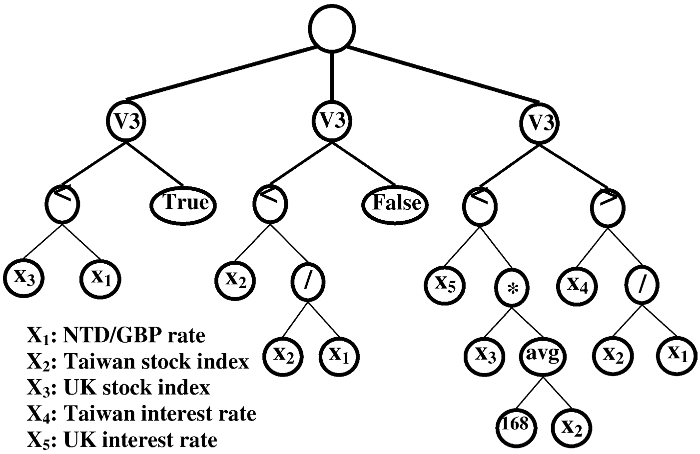

Fig. 10. An Evolved GP Trading Strategy.

The right tree can be interpreted as:

> When Taiwan stock market bulls (suggested by the high Taiwan Stock index), move the funds from UK stock market to UK currency.

The two stock market trading rules recommend allocating funds toward that with indication of higher returns. It advises cashing stocks as a preparation of purchasing foreign stocks, when the foreign stock index looks promising.

9.2. ADF-GP Trading Strategies

ADFs were incorporated for GP to identify possible regularity in profitable trading strategies. However, ADF-GP results are not better than vanilla GP results. We examined those strategies where ADFs were created and evolved by GP and found that most of the ADFs have a constant value of either "*True*" or "*False*." In other words, they are not functions but serve as constants in the trading rules. It is not surprising that this implementation of ADFs give similar returns as the vanilla GP does. As mentioned in Section 5, the time series have been transformed by dividing the daily data with a 250-day moving average. This result indicates that either there is no regularity in the *change of trend* that provides profitable trading or GP is not able to identify such regularity.

Provided with ADFs that are initiated with commonly used technical trading rules, GP still cannot find strategies that give better returns. This indicates that those technical trading rules (and their variations) are not effective on the transformed time series data.

This raises a question about whether the data transformation method used is appropriate for this modeling task or not. In Nikolaev and Iba (2002), they have reported that GP gives different results when the time series are normalized using different data transformation methods. Kassicieh et al. (1998) also reported a similar result when using a GA to make investment decisions. We have normalized the time series by dividing the daily data with 250-day moving average. With such data, GP is searching for patterns in the *change of trend* that give profitable trading. In other words, GP rules can exploit patterns in financial market indices, just like commonly used moving average and filter rules do. When ADFs are incorporated, GP becomes capable of exploiting higher-order complexity, i.e. an ADF gives the first-order pattern while the GP main program calling such ADF defines higher-order complexity (Li & Vitanyi, 1997). We are not certain if higher-order complexity exists in the *change of trend* financial time series. The ADF-GP results do not support this proposition. However, this does not preclude the possibility that such complexity can exist in time series that are normalized using different methods.

We have compared our approach with another work using ADF to find trading strategies in S&P stock index markets and found the author used a different data transformation method in his work: stock indices are divided by 100 while interest rates are divided by 10,000 (Wang, 2000). Similar to our ADF-GP results, Wang's ADF-GP did not discover trading strategies that out-perform B&H in S&P500 spot and future markets. However, his work did not acknowledge ADF-GP is capable of identifying higher-order complexity in the time series. Nor did it mention about the evolved ADF-GP strategies exhibit such complexity. We are inclined to believe that there exist patterns in profitable trading strategies when the time series are applied with appropriate data transformation method. We are currently investigating this hypothesis.

10. CONCLUDING REMARKS

The hot money issue occurred in China has triggered our interest in modeling short-term capital flow in international financial markets. If it is possible to predict such capital inflow and outflow, appropriate measures can be imposed before hand to stabilize global economy. Unfortunately, our finding using GP to simulate a simplified international markets model indicates that such task cannot be accomplished. The devised GP trading strategies do not consistently generate better returns than the buy-and-hold strategy, suggesting that they do not have the ability to predict capital inflow and outflow. Many of the GP strategies recommend the buy-and-hold approach on one or two assets. This indicates that the international short-term capital markets are reasonably efficient, a finding which is similar to that reported by Allen and Karjalainen (1999) and Wang (2000).

However, many GP strategies are able to forecast Taiwan stock market down time and avoid making futile investments. This indicates that GP has the ability to learn from historical data to make profitable trading decisions. Moreover, during market down time when buy-and-hold gives poor returns, many GP strategies are able to identify opportunities and produce better returns than buy-and-hold.

Our investigation of ADF-GP trading strategies does not support the proposition that profitable strategies contain regularity. Nor does it endorse the idea that commonly used technical trading rules are effective on the *change of trend* time series. This seems to counter our intuitions since it is not uncommon for the real-world technical traders apply a combination of technical trading rules to make trading decisions. We are puzzled by this result and have started looking into reasons that have led to such a conclusion. One issue we have identified is the transformation of time series which might have changed the modeling space and time series correlation. Another aspect is the existence of higher-order complexity in financial time series that can be captured by ADF-GP. We are currently investigating different modular GP techniques, in addition to ADFs, and different data normalization methods in order to improve our understanding of regularity in profitable trading strategies.

NOTES

1. Of course, as long as the central government doesn't intervene in the foreign exchange market correspondingly, there are always the direct balance of payments effects, be the foreign investment speculative or not.

2. We have also experimented with the setup where each transaction is executed once in each time step. The preliminary results, however, show very little differences from that of executing a transaction 10 times in a time step. This suggests that for these time series data, trading strategies are not sensitive to the amount of capital flow. In other words, under such time series, a trading strategy gives similar return regardless of the amount of fund transferred in each time step.

ACKNOWLEDGMENTS

An earlier version of the paper has been presented at the 2003 International Conference on Artificial Intelligence (IC-AI03), the 9th International Conference on Computing in Economics and Finance (CEF 2003), and the 3rd International Workshop on Computational Intelligence in Economics and Finance (CIEF'2003).

The second author is grateful for the research support from NSC grant No. NSC 90-2415-H-004-018. We would like to thank members of AI-ECON Research Center, particularly Chueh-Yung Tsao and Bin-Tzong Chie, for spending time to discuss financial trading with us during the writing of this paper. We also thank the reviewers for their comments and suggestions and Ingrid Peterson for proof reading this paper.

REFERENCES

Allen, F., & Karjalainen, R. (1999). Using genetic algorithms to find technical trading rules. *Journal of Financial Economics*, *51*(2), 245–271.

Baba, N., Inoue, N., & Asakawa, H. (2000). Utilization of neural networks & GAs for constructing reliable decision support systems to deal stocks. Working Paper of Information Science, Osaka-Kyoiku University, Asahiga-Oka, 4–698–1, Kashiwara City, Osaka Prefecture, Japan.

Brock, W., Lakonishok, J., & LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, *47*, 1731–1764.

Chen, S.-H., & Kuo, T.-W. (2002). Genetic programming: A tutorial with the software simple GP. In: S.-H. Chen (Ed.), *Genetic Algorithms and Genetic Programming in Computational Finance* (pp. 55–77). Kluwer Academic Publishers.

Chen, S.-H., & Kuo, T.-W. (2003). Discovering hidden patterns with genetic programming. In: S.-H. Chen & P. P. Wang (Eds), *Computational Intelligence in Economics and Finance* (pp. 329–347). Springer-Verlag.

Kassicieh, S. K., Paez, T. L., & Vora, G. (1997). Investment decisions using genetic algorithms. In: *Proceedings of the Thirtieth Hawaii International Conference on System Sciences* (pp. 484–490).

Kassicieh, S. K., Paez, T. L., & Vora, G. (1998). Data transformation methods for genetic-algorithm-based investment decisions. In: *Proceedings of the Thirty-First Hawaii International Conference on System Sciences* (pp. 122–127).

Koza, J. (1992). *Genetic programming: On the programming of computers by means of natural selection*. Cambridge, MA: MIT Press.

Koza, J. (1994). *Genetic programming II: Automatic discovery of reusable programs*. Cambridge, MA: MIT Press.

Li, M., & Vitanyi, P. (1997). *An introduction to Kolmogorov complexity and its applications*. Springer Verlag.

Neely, C., & Weller, P. (1999). Technical trading rules in the European monetary system. *Journal of International Money and Finance*, *18*, 429–458.

Neely, C., Weller, P., & Dittmar, R. (1997). Is technical analysis in the foreign exchange market profitable? A genetic programming approach. *Journal of Financial and Quantitative Analysis*, *32*(4), 405–426.

Nikolaev, N., & Iba, H. (2002). Genetic programming of polynomial models for financial forecasting. In: S.-H. Chen (Ed.), *Genetic Algorithms and Genetic Programming in Computational Finance* (pp. 103–123). Kluwer Academic Publishers.

Pesaran, M., & Timmermann, A. (1995). Predictability of stock returns: Robustness and economic significance. *Journal of Finance*, *50*, 1201–1228.

Pring, M. J. (1991). *Technical analysis explained* (2nd ed.). New York: McGraw-Hill.

Seshadri, M. (2003). *Comprehensibility, overfitting and co-evolution in genetic programming for technical trading rules*. Master's Thesis. Worcester Polytechnic Institute.

Syswerda, G. (1991). A study of reproduction in generational and steady-state genetic algorithms. In: G. J. E. Rawlins (Ed.), *Foundations of Genetic Algorithms* (pp. 94–112). Morgan Kaufmann.

Wang, J. (2000). Trading and hedging in S&P 500 spot and futures markets using genetic programming. *Journal of Futures Markets, 20*(10), 911–942.

Yu, T., & Bentley, P. (1998). Methods to evolve legal phenotypes. In: *Proceedings of the Fifth International Conference on Parallel Problem Solving From Nature* (pp. 911–942). Springer.

Uncited references

References cited in the text must appear in the reference list; conversely, each entry in the reference list must be cited in the text... The author must make certain that each source referenced appears in both places and that the text citation and reference list entry are identical in spelling and year.

Chen and Kuo (2003) and Neely and Weller (1999).