

A Game-Theoretical Approach for Designing Market Trading Strategies

Garrison W. Greenwood and Richard Tymerski

Abstract— Investors are always looking for good stock market trading strategies to maximize their profit. Under the technical school of thought trading rules are developed by studying historical market data to find trends that investors can exploit. These market trends tend to appear when certain features (narrow range, DOJI, etc.) appear in the historical data. Unfortunately, these features often appear only in partial form, which makes trend analysis challenging.

In the paper we co-evolve fuzzy trading rules from market trend features. We show how fuzzy membership functions naturally handle partial form features in historical data. The co-evolutionary process is formulated as a zero-sum, competitive game to match how trading strategies are evaluated by brokerage firms. Our experimental results indicate the co-evolutionary process creates trading rule-bases that produce positive returns when evaluated using actual stock market data.

I. INTRODUCTION

The recent emergence of online trading has made the stock market accessible to small investors. Brokerage firms work on behalf of investors to purchase quantities of a single stock. (Investors pay a small fee for every transaction.) Some discount brokerage firms provide little or no investment advice, which means it is up to the small investor to come up with their own investment strategies. Other brokerage firms do advise investors when to buy or sell stock, but the underlying strategy is proprietary and for that reason is not disclosed to outsiders. Strategies are continually refined because markets can be volatile and good strategies that produce high returns tend to attract more investment dollars—and higher commissions for the brokers!

There are two schools of thought on developing investment strategies. In the *fundamental approach* decisions about buying, selling or holding a company's stock arise from a careful analysis of the the company's books. On the other hand, in the *technical approach* studying a company's past trading activity can help predict future stock prices. In this paper we focus on the technical approach.

For decades people have proposed various technical trading rules but the majority of this literature has found that, for the most part, the investment rules just don't work. Indeed, one researcher [1] even went so far as to dismiss technical investment rule methods entirely! But the advent of more powerful and inexpensive computers has promoted new and powerful data mining methods that can search high frequency market data for trading activity patterns. The reported demise of the technical school of thought was clearly premature.

G. Greenwood and R. Tymerski are both with the Electrical & Computer Engineering Department at Portland State University, Portland, OR 97207-0751 USA (email: greenwood, tymerski@ece.pdx.edu)

Computational intelligence (CI) offers a variety of useful techniques for constructing trading rules via technical approaches. For example, Allen and Karjalainen [2] created trading rules using genetic programming. The goal was to develop a function that returns either a 'buy' or 'sell' signal. (Decisions to 'hold' a stock were not implemented.) The difference between an excess return¹ and a simple buy-and-hold strategy² over a finite training period measures the fitness of a rule. The authors claim their evolved rules had reduced trading volatility, but their rules do not lead to higher absolute returns than a buy-and-hold strategy. Not to be deterred, the authors claim some investors may still find their rule-base worthwhile.

Potvin et. al [3] also tried genetic programming to evolve investment rules. Their method was computationally expensive and tended to converge prematurely (no reported improvements after 50 generations). They too showed poorer performance than a buy-and-hold approach, especially in a rising market. The authors claim their trading rules were "generally beneficial", but only when the market is stable or falling.

Dourra and Siy [4] used fuzzy logic to create a trading strategy. The system inputs were momentum indicators, such as the rate of change of stock prices over a defined period. Gaussian membership functions and a fuzzy rule-base, hand-crafted from expert's knowledge completed the system. A Mamdani fuzzy implication method output a single value between 0 (strong sell) and 100 (strong buy). Their fuzzy system was evaluated on 3 years of stock prices from several different companies. The investment returns, based on buy or sell recommendations from the fuzzy system, were exceptionally high—some yielding over 250% returns! However, those returns should come as no real surprise since the membership functions were deterministically adjusted to fit the training data. Consequently, the membership functions used for the Intel Corporation stock were different from those used for the General Motors stock. Moreover, the rule antecedents were very sophisticated—some involved both conjunctions and disjunctions of fuzzy variables—but the authors did not say how those rules were derived.

Lam [5] also constructed a fuzzy trading rule-base, but the rule-base was evolved with a genetic algorithm that selected a subset of rules from a set of 36 pre-defined fuzzy

¹A risk-free return is the return of an asset, such as T-bills, with no risk whatsoever. An excess return is the return received over and above a risk-free return.

²In this strategy the investor buys a stock and then holds it for a long time period, hoping to outlast any market fluctuations.

rules. The antecedent of each rule was a conjunction of several moving average trends such as daily moving average, weighted moving average and exponential moving average. Other fuzzy rules had antecedents with momentum terms such as relative strength index, rate-of-change and fast and slow stochastic. The output (consequent) of each rule was a buy, sell or hold order using the Sugeno method of fuzzy inference. The system outperformed a buy-and-hold strategy, but it is not clear how the system is trained. The author claims the system was trained for m days and then applied to trade for n more days *before re-training*(?) It was never stated how the fitness for the m training days relates to the fitness for the n trading days. The paper is useful only from the standpoint it shows the potential of evolving fuzzy trading rule-bases.

In this paper we describe some preliminary work using a *coevolutionary algorithm* to evolve a family of fuzzy trading rule-bases, each of which serves as a unique trading strategy. Researchers in the past have evolved just a single population of strategies. What differentiates our work from all previous evolved rule-base work is we coevolve independent populations of strategies under a game-theoretical environment. This approach seems quite natural because of the close relationship between competitive coevolution and game theory. Indeed, evolutionary game theory can help explain the underlying dynamics of coevolutionary algorithms [6]. From a practical standpoint coevolution makes perfect sense since it closely matches how investment firms develop their own internal strategies. This aspect is discussed further in the next section. Experimental results are included to demonstrate our proposed method.

II. WHY USE GAME THEORY TO DEVELOP TRADING STRATEGIES?

Investment counselors, working for brokerage firms, develop trading strategies for stock market investments. Good trading strategies attract more investor dollars while poor strategies discourage further investments. Each firm internally develops not just one, but a set of investment strategies to deal with market volatility. For example, the firm will need one strategy to deal with investments in a bull market, where stock prices generally rise, and another strategy for bear markets, where stock prices generally fall. Investment counselors adapt their individual trading strategies based on how well those strategies perform. The collective strategies from all of the investment counselors comprises the investment services a brokerage firm offers to a potential investor. This process can be envisioned as an evolving set of strategies where fitness is measured by the returns derived by using the strategy to make investment decisions.

But how does a brokerage firm know if an evolved set of strategies is any good? That question illustrates the problem with evolving a single population of trading rules. Intuitively strategies that provide the greatest returns are the most appealing and would be assigned high fitness. But defining fitness proportional to returns is overly simplistic and may

even give a false picture of how good a strategy actually is. An example will help illustrate the problem.

Suppose two brokerage firms independently evolve a set of investment strategies. Assume the best performing strategy from the first brokerage firm consistently yields a 4% return over a 90-day trading period. If fitness is proportional to returns, then this strategy, by definition, has the highest fitness. But does a 4% return really *qualify* as high fitness? The only way to know for certain is to compare that 4% return against what the other brokerage firm can offer. If the best strategy from the second firm only has a 1.5% return, then 4% is pretty good and should signify high fitness. On the other hand if the second firm can offer a 7% return strategy, then a 4% return does not qualify as high fitness.

It is important to differentiate between the relative fitness and the true fitness of an investment strategy. Relative fitness compares returns from strategies within a single brokerage firm whereas true fitness contrasts returns from strategies between independent firms. Investors compare the returns achieved by various brokerage firms before deciding where to invest their money. Hence, true fitness measures the ability to attract investor dollars. Survival, therefore, should depend on true fitness rather than relative fitness. The point here is there is no way to say whether or not a given return constitutes high fitness unless it is compared against the returns from competing brokerage firms. Consequently, strategies should arise from competitive coevolution [7].

Stock market investment is naturally expressed as a game. The players are brokerage firms, which independently develop investment strategies. Strategies compete against each other in the marketplace and receive payoffs in the form of greater or fewer investor dollars depending on how well they perform. (Poor strategies lose investments as investors switch to better performing strategies.) Hence, stock market investment is essentially a zero-sum game with the strategies formed through competitive coevolution.

NOTE: In some games players make decisions based on historical data, which is common knowledge. One of the best examples is the widely studied *minority game* [8]. The stock market game formulated here belongs to this same family of games.

III. PROBLEM FORMULATION

The objective is to design a strategy for trading shares of a single stock. Stock price data, recorded over a large number of consecutive trading days, is available to help develop the strategy. The strategies are based on finding conditions in market historical data that predicts subsequent up-trend days³ Appropriate investments (or positions) are taken only on the open of a predicted trend day and are exited at its close.

Four data items were recorded each day: the open share price (O), the high (H) and low (L) for the day and the

³In this work we did not attempt to find down-trend days, which would require a different set of strategies.

closing share price (C). Two types of trading days are of interest:

Definition: (*up-trend day*)

$$\begin{aligned} O &\leq L + 0.1(H - L) \\ C &\geq H - 0.2(H - L) \end{aligned} \quad (1)$$

Definition: (*down-trend day*)

$$\begin{aligned} O &\geq H - 0.1(H - L) \\ C &\leq L + 0.2(H - L) \end{aligned} \quad (2)$$

Qualitatively, this simply means for an up day the opening price is close to the day's low and the closing price is close to the day's high. Similarly, for a down-trend day the opening is at or near the high and the close is near or at the low for the day. An interesting property of trend days—which keen investors can exploit—is the high-low differential tends to be relatively large.

Studies have identified several trading data features that often proceed up or down-trend days. These features, described below, are defined in terms of O , C , H and L . Today is indexed with $i = 0$ and previous days are indexed $i = 1, 2$, and so on. $O[-1]$ denotes the next trading day opening price.

- NR k

With $H[i]$ and $L[i]$ denoting the high and low for the i -th day, the range is defined as $R[i] = H[i] - L[i]$. NR k exists if today's range is less than the ranges for the previous $k - 1$ days. That is,

$$R[0] < \min(R[1], \dots, R[k - 1]) \quad (3)$$

NR k days represent volatility contraction, which often-times leads to volatility expansion in the form of wide range days. The greater the number of narrow range days, the greater the counter reaction in wide ranging days.

- DOJI

DOJI indicates that the open and close for the trading day are within some small percentage (x) of each other. A DOJI means the market reflects temporary price indecision and often signals a major reversal in the market. DOJI is a predicate function—i.e., it returns 1 (TRUE) or 0 (FALSE). It is defined as

$$\text{DOJI}(x) = \begin{cases} 1 & |O - C| \leq x \cdot (H - L) \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

- Hook day

A hook day occurs when the price opens outside the previous day's range and then proceeds to reverse direction, generally indicating a reaction to temporarily overbought or oversold market conditions.

There are two versions of a hook day. For the up hook day $O[-1] < L[0] - \delta$ and for the down hook day

$O[-1] > H[0] + \delta$. In both cases $0 \leq \delta \leq 10$ is user selectable.

IV. FUZZY SYSTEM DESIGN

Given stock market historical data, it is always possible to (deterministically) analyze it and mark where any of the features described in the previous section are present. But a more subtle and challenging problem must be dealt with. The rule-base contains a set of fuzzy rules that predict whether the next trading day is likely to be a trend day. Investment decisions—i.e., whether to buy, sell or hold—are made based on these predictions. Fuzzy rules are used because crisp rules are too restrictive. An example will help fix ideas.

The crisp rule “**if** NR7 **then** ...” is true if and only if the range during the current day $R[0]$ is less than the range during any of the previous six days $R[1], R[2] \dots R[6]$. The rule won't be true if even one of the previous ranges is less than $R[0]$. But suppose the inequality is satisfied for say five out of the six days, which makes the antecedent *almost* true? As another example, Nison [9] asked

“How do you decide whether a near-DOJI day (that is, where the open and close are very close, but not exact) should be considered a DOJI? This is subjective and there are no rigid rules...”

Fuzzy reasoning can effectively deal with such uncertainties. Crisp rules, which must give either a 'yes' or a 'no' answer, cannot handle these situations but fuzzy rules can because they can also provide fuzzy answers somewhere in between a 'yes' or 'no'.

In our approach stock market data is analyzed to determine how closely it matches the formal definitions of the features described in the previous section. Membership functions return a value between 0 and 1 indicating to what degree features are present. The resultant fuzzy variables are then collected into fuzzy if-then rules, which constitutes the trading rulebase. The outputs of active rules—i.e., rules whose antecedent are satisfied—are combined into a fuzzy output variable. This variable is defuzzified to produce a crisp value on the unit interval, which the desirability of buying stock on the next trading day.

A. Membership Functions

Fuzzification is the process that maps days (D) onto the unit interval via a membership function $\mu(D)$. More precisely, D represents the number of previous days that a particular feature is satisfied. For instance, for the NR7 feature $D \in \{0, 1, \dots, 6\}$. Then $\mu(0) = 0$ means the NR7 definition was not satisfied during any of the six previous days, $\mu(6) = 1$ means the definition was satisfied during all six previous days (i.e., NR7 is definitely present) and $0 < \mu(D) < 1$ means NR7 was satisfied for some $D < 6$ days. Trapezoidal membership functions are most appropriate for the most of the features (see Figure 1).

- NR k

For this feature the equation is slightly different for each value of k . Let $\mu_k(x)$ denote the membership function for NR k . Then

$$\mu_k(x) = \begin{cases} 0 & x < v_{\min} \\ c(x - v_{\min}) & v_{\min} \leq x < v_{\max} \\ 1 & x \geq v_{\max} \end{cases} \quad (5)$$

with parameter values as shown in Table I.

k	c	v_{\min}	v_{\max}
4	1/2	2	4
6	1/3	3	6
7	1/3	4	7

TABLE I

PARAMETER VALUES FOR NARROW RANGE FEATURES

In the above equation $x = D + \eta$ where $D \leq k$ is the number of days where (3) holds and, with $\tilde{R} = \max_{1 \leq j \leq k} R[j]$,

$$\eta = \frac{\tilde{R} - R[0]}{\tilde{R}}$$

Notice that η increases the membership value for smaller previous day ranges.

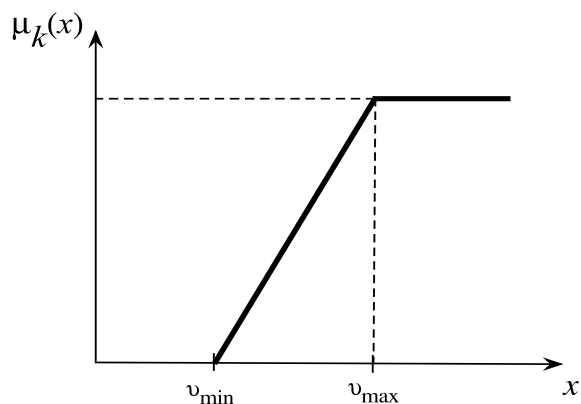


Fig. 1. Membership functions for the features.

- DOJI

For this feature the membership function equation is

$$\mu(x) = \begin{cases} 1 - x/\rho & 0 \leq x \leq \rho \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where typically $\rho \in [0.05, 0.30)$. x represents the percent difference between O and C and ρ represents the threshold percentage.

- Hook Day

The formula for this membership function is

$$\mu(x) = \begin{cases} 0 & x < -\frac{1}{2} \\ 2(x + 0.5) & -\frac{1}{2} \leq x < 0 \\ 1 & x \geq 0 \end{cases} \quad (7)$$




Fig. 2. DOJI membership function with $\rho = 0.1$.

where $x = L[0] - \delta - O[-1]$ for an up hook day and $x = O[-1] - \delta - H[0]$ for a down hook day.




Fig. 3. Hook day membership function

B. The Fuzzy Rule-Base

Unfortunately, just detecting the presence or absence of a single feature is not a very good trend day predictor. The problem is to find combinations of features that make a good trend day predictor⁴.

If there are N total features, then there are N total rules in the rule-base⁵. Consider the rule

if x is NR4 **then** output is up-trend day

The semantics of this rule is as follows. The term “ x is NR4” means ranges for the current and the previous three days are computed. x represents how many of those days meet the NR4 definition. This crisp data value is the argument for the NR4 membership functions which returns

⁴Certain real parameter values must also be chosen. For example, the percentage threshold input value is needed for the DOJI feature.

⁵There would be $2N$ total rules if both up-trend and down-trend days are predicted.

a number between 0 and 1 to give the degree of membership for NR4. The output value is the degree of membership for an up-trend day.

A singleton fuzzifier is used for the inputs. The set of possible outputs is

$$\Lambda = \{ 0.25 \quad 0.5 \quad 0.75 \quad 1.0 \}$$

These output values represent the likelihood an investor would buy stock because a given fuzzy rule had fired.

The fuzzy rule-base is encoded as a matrix M with one column for every $\lambda_i \in \Lambda$ and one row for every fuzzy rule [10]. Each rule is of the form “if x is feature then y is an up-trend day”. In this work we investigated 5 features so M has 5 rows corresponding to, from top to bottom, the features NR4, NR6, NR7, DOJI, and Up Hook Day (with $\delta = 0.5$). There are four columns, one for each value in Λ . A typical matrix might look like

$$M = \begin{pmatrix} 0 & 0.6 & 0.5 & 0 \\ 0 & 0.33 & 0.33 & 0.33 \\ 0 & 0.1 & 1.0 & 0.9 \\ 0 & 0.44 & 0.5 & 0.1 \\ 0 & 0.1 & 0.2 & 0.7 \end{pmatrix}$$

where each $m_{ij} \in M$ is a weight. The first row in the above matrix is thus interpreted as the investor saying *If the current day is the feature NR4, then I think the likelihood of an up-trend day tomorrow is 0.5 or 0.75 (out of 1) with strengths 0.6 and 0.5, respectively*. Notice the strengths do not have to sum to 1.0.

Rule processing is straightforward. Given the rule-base matrix M , the vector of fuzzy numbers

$$A = (\mu_{NR4}, \mu_{NR6}, \mu_{NR7}, \mu_{DOJI}, \mu_{hook})$$

is created by running the training data through the individual membership functions. A new A vector is created for each trading day. For example, on the 50th training day the ranges $R(47)$, $R(48)$ and $R(49)$ are computed and $\mu_4(x)$ is computed using (5). Similarly the other membership function values are determined using (6) and (7). Then $A \circ M = B$ where

$$b_j = \max_{1 \leq i \leq 5} \min \{a_i, m_{ij}\}$$

The fuzzy output vector B is then defuzzified to get a crisp output value indicating the desirability of buying shares of stock. A center of average defuzzification was used. That is,

$$U = \frac{\sum_{i=1}^N b_i \cdot \lambda_i}{\sum_{i=1}^N \lambda_i} \quad (8)$$

where λ_i is the i -th singleton from Λ and N is the number of active rules.

The crisp output $U \in [0, 1]$ indicates the likelihood of an up-trend day or, equivalently, the desirability of purchasing more stock shares. However, stock is only purchased if the desirability exceeds a user-selected threshold. We chose 0.8

for this threshold. If the output value is greater than 0.8, then a ‘buy’ signal is generated and how high above 0.8 determines the number of shares to buy. Specifically, the amount of stock purchased increased linearly the higher the output was above 0.8, subject to sufficient funds in the bank account, up to a maximum of 20 shares. Stock was bought at the opening price and sold at the closing price. All proceeds were deposited into the bank account at the end of each trading day.

C. The Coevolutionary Algorithm

The fuzzy rule-base was evolved using a (40 + 40) evolution strategy. Both recombination and mutation are used as variation operators. Each offspring are created by randomly choosing two parents. After appending together the columns of the M matrix in each parent to form a real vector, standard 1-point crossover, followed by mutation, creates the offspring. Mutation added a gaussian distributed random variable with zero mean and standard deviation $\sigma = 0.2$ to each term in the M matrix. The mutation was redone if the component value was less than 0.1 or more than 1.0. For the defuzzified output any value below 0.5 was automatically set to 0 because a desirability of buying stock with a value that low makes no sense. The evolutionary algorithm was run for 750 generations.

Both the α and the ω rule-bases—corresponding to the α and the ω brokerage firms—evolve independently. Fitness of a rule-base is determined by tournaments with the competitor rule-bases. Specifically, each individual in the α population was provided with \$400 and then used its rule-base as an investment strategy over 150 consecutive trading days with a start date chosen randomly from the first 850 trading days in the training market database. The ω population did the same over the same 150 trading days. Afterwards each individual in both populations has an ending bank account balance after selling all shares purchased during the trading period. The bank account balance of each individual in the α (ω) population was compared against 10 randomly chosen individuals from the ω (α) population. The individual with the largest bank account records a win.

Winning in this game has a positive payoff while losing has a negative payoff. Every time an individual wins against an individual in its tournament set, the winner gets 2% of the loser’s bank account balance. Conversely, if the individual loses against a member of its tournament set, he has to pay 2% of his bank account balance to winner. Hence, individuals get paid or have to pay every time they participate in a tournament. (Accounts were not actually credited or debited until after all tournaments were finished.)

This payoff approach to conducting tournaments emulates real-world trading. Brokerage firms compete for investor dollars. The bank account associated with each individual can be thought of as money given to a brokerage firm by investors. Each firm starts out with the same amount and trades according to the strategy defined by its fuzzy rule-base. Firms with lower bank accounts at the end of the trading period have poorer strategies and poor strategies

cause investors to move their money elsewhere—i.e., to better performing firms.

V. EXPERIMENTAL RESULTS

A database of 1600 trading days were available for training and testing. The first 1000 days are reserved for training whereas the 600 remaining days are used for validation testing. Each training epoch consists of 150 consecutive trading days, with a random starting date and partitioned into three overlapping 90 day segments. The second segment used the last 60 days from the first segment plus 30 new days; the third segment used the last 30 days of the first segment plus 60 new days. The overlaps help mitigate market volatility.

Each training epoch started with a \$400 bank account. The bank balance at the end of the first 90 day training segment was the beginning bank balance of the second 90 day segment. Similarly the third segment beginning balance was the same as the second segment ending balance. (Bank account balances are adjusted as explained in the previous section.) Testing was conducted in the same way, except the 150-day window was chosen from the latter 600 days in the database.

Figure 4 shows the results from a typical run on a 90-day trading session from the testing portion of the database. Notice the desirability exceeded the 0.8 threshold on only a few days.

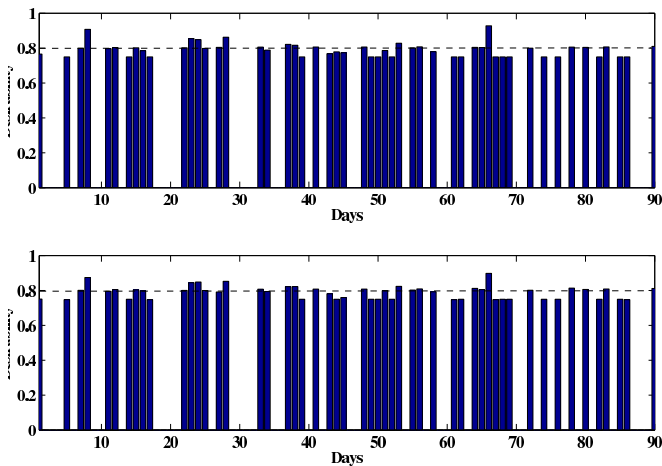


Fig. 4. Typical trading behavior for the best fuzzy trading rule sets over a 90 trading period. Shares were bought whenever the desirability exceeded 0.8.

Table II shows the bank account balances for the top five and bottom five trading strategies after completing the 150 days of trading using the training database. Notice for both firms the top five strategies had positive returns whereas the bottom five strategies had negative returns. The positive returns are considerably higher than the starting balance of \$400. Remember the ending bank account balances are not due solely to buying stock at a reasonable price. Good performing strategies have a higher payoff by taking money out of the bank accounts of poorer performing strategies. This zero-sum game format matches how investors in the

real world choose brokerage firms: firms with good strategies attract more investor's dollars while firms with poorer strategies drive away investor's dollars.

	Account Balances from Top 5 Strategies (\$)	Account Balances from Bottom 5 Strategies (\$)
α Firm	548.4, 498.2, 489.6, 487.8, 479.6	336.1, 335.4, 332.2, 331.1, 326.2
ω Firm	548.9, 548.7, 548.3, 547.6, 547.0	322.2, 321.2, 319.0, 304.5, 300.1

TABLE II

BANK ACCOUNT BALANCES FROM TOP AND BOTTOM TRADING STRATEGIES FROM TWO BROKERAGE FIRMS. RESULTS ARE OBTAINED OVER A 150 TRADING PERIOD WITH AN INITIAL \$400 BANK BALANCE.

The best α firm fuzzy rule-based strategy was then compared against a simple investment strategy over 90-day trading periods using the test market data—i.e., the last 600 days in the market database, which were not used for training. In the simple investment strategy fifty shares of stock are purchased at the beginning of the trading day and then sold at the closing price at the end of the trading day. Two trading periods were chosen. The first trading period begins on day 1200 while the second trading period begins on day 1400. Figures 5 and 6 compares the simple and fuzzy trading strategies for the first and second trading periods, respectively.

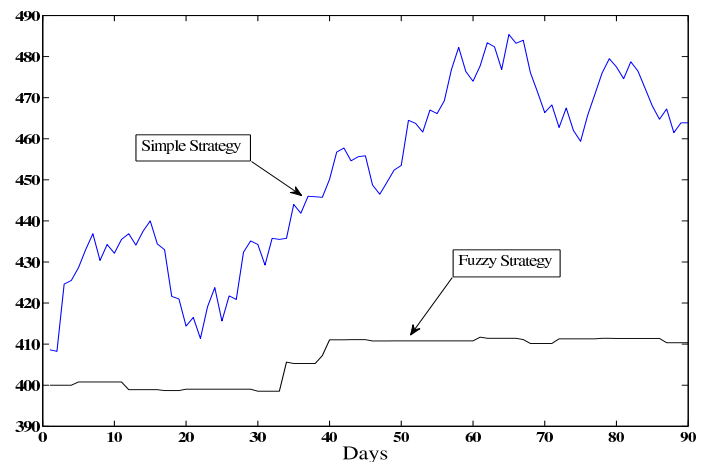


Fig. 5. Bank account balances for the top α firm strategy and the simple investment strategy over a 90 day period of test market data starting at the 1200th trading day. Both accounts started with a \$400 balance.

Some interesting observations emerge from a close examination of Figures 5 and 6. The returns from buying and selling stock every day indicates how the market is doing. The market data shows high volatility. Notice the market is generally up for several weeks after the 1200th trading day but reverses dramatically after the 1400th trading day. More to the point, the simple investment strategy had a +15.7% return in the first trading period but a -8.2% return during the second trading period. Clearly buying stock every day is not a good investment strategy. Conversely, the fuzzy rule-based

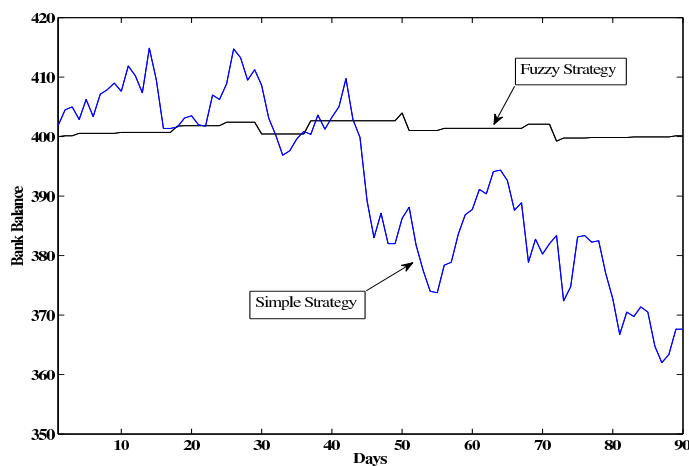


Fig. 6. Bank account balances for the top α firm strategy and the simple investment strategy over a 90 day period of test market data starting at the 1400th trading day. Both accounts started with a \$400 balance.

had a modest +2.5% return during the first trading period but no loss during the second trading period. Although the returns were not as great when the market was up, it did protect capital when the market was down.

VI. FUTURE WORK

If the two brokerage firms only used their best trading strategy, in principle one firm could do some reverse engineering and ultimately deduce the other firms trading strategy (i.e., the fuzzy rules). Notice in Table II that the returns between any of the top two strategies are not significantly different. Hence, a mixed strategy will provide security without sacrificing much of the return.

Both brokerage firms used the same membership functions. In practice each brokerage firm would use its own internally developed membership functions. In this work the membership functions were fixed but in our future experiments we will evolve the membership functions as well.

Our evolved trading rule-base only bought stock when the market data suggested a pending up-trend day. We need to also evolve a set of down-trend day trading rules. The up-trend day rules would indicate buying opportunities while the down-trend days would indicate selling short to help minimize losses.

Finally, our trading strategy relies only on opening and closing stock prices. Day traders make multiple trades throughout the day to maximize returns. Figure 7 shows the price variation during a 6-hour trading period.




Fig. 7. Stock prices over a single trading day period.

Notice that buying the stock at the beginning of the trading session and selling the stock after 4 hours would produce a higher profit than selling or holding based on the opening and closing prices alone, which is what our fuzzy rule based decisions on. This suggests fuzzy rules for intra-day trading would produce even higher profits. We will investigate this issue as well.

REFERENCES

- [1] E. Fama. Efficient capital markets: a review of theory and empirical work. *Journal of Finance*, 25:383–417, 1970.
- [2] F. Allen and R. Karjalainen. Using genetic algorithms to find technical trading rules. *Journal of Financial Economics*, 51(2):245–271, 1999.
- [3] J. Potvin, P. Soriano, and M. Vallee. Generating trading rules on the stock markets with genetic programming. *Computers & Operations Research*, 31:1033–1047, 2004.
- [4] H. Dourra and P. Siy. Investment using technical analysis and fuzzy logic. *Fuzzy Sets & Systems*, 127:221–240, 2002.
- [5] S. Lam. A genetic fuzzy expert system for stock market timing. In *Proceedings of the 2001 IEEE Congress on Evolutionary Computation*, pages 410–417, 2001.
- [6] S. Ficici and J. Pollack. A game-theoretic approach to the simple coevolutionary algorithm. In *Proceedings of the Sixth International Conference on Parallel Problem Solving from Nature*, pages 467–476, 2000.
- [7] C. Rosin and R. Belew. New methods for competitive coevolution. *Evolutionary Computation*, 5(1):1–29, 1997.
- [8] T. Lo, P. Hui, and N. Johnson. Theory of the evolutionary minority game. *Physical Review E*, 62(3):4393–4396, 2000.
- [9] S. Nison. *Japanese Candlestick Charting Techniques: A Contemporary Guide to the Ancient Investment Techniques of the Far East*, page 150. New York Institute of Finance, 1991.
- [10] J. West and B. Linster. The evolution of fuzzy rules as strategies in two-player games. *Southern Economic Journal*, 69(3):705–717, 2003.