Journal of Management and Economics
http://www.econ.uba.ar/www/servicios/publicaciones/journal3/index.htm
November 2000

EDDIE In Financial Decision Making

Edward P.K. TSANG[1], Jin LI[1], Sheri MARKOSE[2], Hakan ER[2], Abdel SALHI[3], Giulia IORI[4]
[1] *Department of Computer Science*
[2] *Department of Economics*
[3] *Department of Mathematics*
[4] *Department of Accounting, Finance and Management*
University of Essex, Wivenhoe Park, Colchester, United Kingdom
http://cswww.essex.ac.uk/CSP/finance

SUMMARY

This paper gives an overview of the EDDIE project. It describes the principles and applications of EDDIE in making financial decisions, including applications to share prices and indices forecasting and arbitrage. EDDIE is designed as an interactive decision tool, not a replacement of expert knowledge. Experts channel their knowledge into the system through (a) selection and preparation of data and (b) providing feedback to EDDIE. EDDIE's main role is to explore interactions between variables and to find thresholds for the variables. Performance of EDDIE depends on both the quality of the users' input and the efficiency of its genetic programming based search engine.

KEY WORDS: financial forecasting; genetic programming

* Contact author:
Professor Edward Tsang, Department of Computer Science, University of Essex, Colchester CO4 3SQ, UK; tel: +44 1206 872774; fax: +44 1206 872788; email:edward@essex.ac.uk

1. Introduction – New Opportunities From New Technology

Can financial markets be predicted? According to the Efficient Market Hypothesis (EMH), the answer is no [Malkiel 1992]. However, there have been many questions over the EMH. Here let us focus on the computation aspect of EMH. For a market to be efficient, information must be fully discounted by investors. The question is: how well can an investor process and act on new information? Does *consequential closure* hold in human reasoning?[1] The answer must be no, as human beings have limited computation capability, even with the help of computers (if consequential closure holds, chess will be no more interesting than tic-tac-toe). That means we cannot possibly make all the inferences from all the information available to us. All being equal, investors who are capable of making more inferences will have an edge over those who cannot. Faster computers and more efficient algorithms allow us to make far more inferences than before, hence they bring new opportunities. This is where our research comes in.

Many factors could directly or indirectly affect the future price of an investment. Such factors are often inter-related, which adds to the difficulty of analysis. The combinatorial explosion problem prevents one from examining combinations of all the factors and all possible interactions between them. EDDIE (which stands for *Evolutionary Dynamic Data Investment Evaluator*) is an interactive tool, designed at University of Essex, to help analysts to search the space of interactions and make financial decisions. This paper describes the principles of EDDIE and its applications to share price and index forecasting and arbitrage.

2. EDDIE As A Decision Making Tool

In this section, we introduce the basic ideas of EDDIE, and explain its potential role for financial decision making. Figure 1 shows the way that EDDIE can be used as a financial forecasting and decision making tool. It should be emphasised that EDDIE does not replace the role of experts. It serves to improve the productivity of users who may have various levels of expertise in the domain.

[1] If *consequential closure* holds for the *know* relation, that means if one knows "P implies Q" then whenever one knows P, one must also know Q.

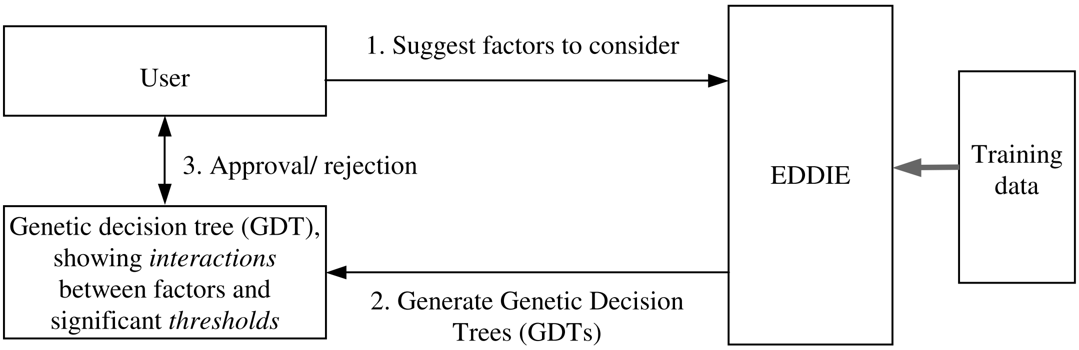

Figure 1. The role of EDDIE as a financial forecasting and decision making tool

One class of predictions that EDDIE can make takes the form:

"Will the price of X rise (or fall) by *r%* within the next *n* days?"

EDDIE is designed to work with financial experts. To search for rules for predicting investment opportunities, the users are responsible for suggesting a set of factors which they consider relevant. This is a point where expert knowledge can be channelled into the computer program. For example, the user might consider closing prices, interest rates, or money supply as being relevant to the prediction of share price movements. These factors may interact with each other. The space of models explaining the interaction between these factors is huge, and growing exponentially with the number of factors given. For example, it may be the combination of money supply and the interest rate that affects the movement of a particular index, not these two factors independently. Besides, the space of thresholds is huge too. For example, whether the price-earning ratio of a particular share is above 16.5 may be critical to the prediction of its price movement. The role of EDDIE is to help the user to search in the huge space of plausible models to find those which provides a good explanation to past data (which is often referred to as *training data*).

EDDIE generates *Genetic Decision Trees* (GDTs, to be explained later), which the human user may approve (rules that agree with the user's expert knowledge) or reject (rules which extract historical patterns that the user considers irrelevant to the future). The user may also ask EDDIE to build alternative GDTs based on a modified set of factors. This generate-and-approve/reject cycle continues until the user is satisfied with the GDTs generated or is convinced that no useful patterns can be extracted from the historical data by EDDIE.

EDDIE searches the space of models using genetic programming [Koza 1992], which has been successful in many applications, including financial applications (e.g. see [Angeline & Kinnear 1996; Koza et al 1996; Neely et al 1997; Butler 1997; Chen 1997; Kaboudan 2000]). In EDDIE, a candidate solution is represented by a *genetic decision tree* (GDT). The basic elements of GDTs are *rules* and *forecast values*, which correspond to the *functions* and *terminals* in GP. Figure 2 shows an example of a simple GDT. A useful GDT in the real world is almost certainly a lot more sophisticated than this. In GP terms, the questions in the example GDT are *functions*, and the proposed actions are *terminals*, which may also be forecast values. In this example, the GDT is binary; in general, this need not be the case.

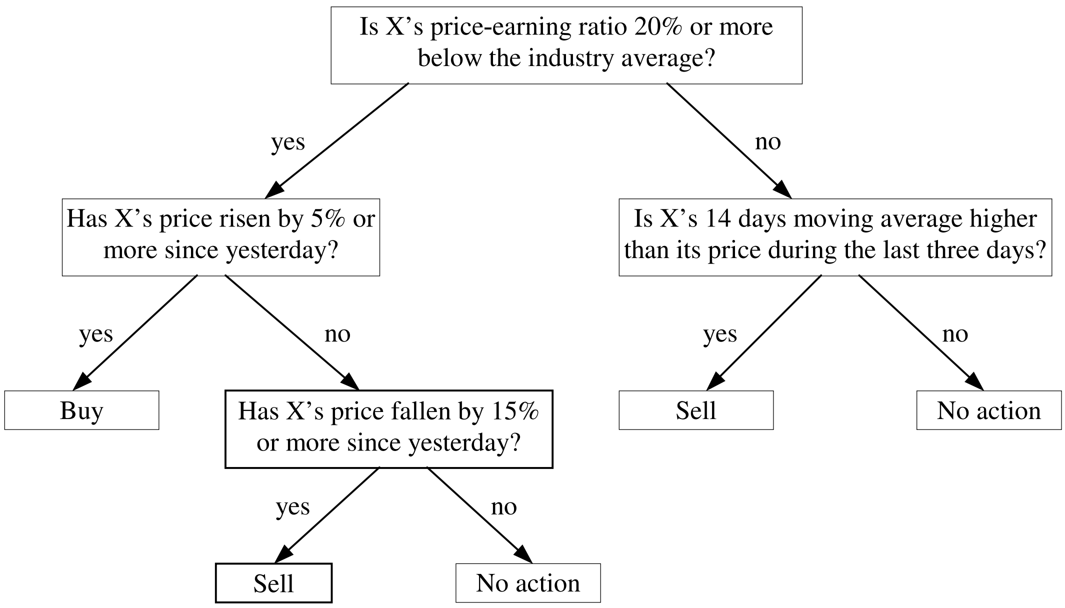

Figure 2. A (simplistic) GDT concerning the actions to take with Share X

A GDT can be seen as a set of rules. For example, one of the rules expressed in the GDT in figure 2 is:

IF X's price-earning ratio is 20% or more below the industry average

AND X's price has risen by 5% or more since yesterday,

THEN Buy X.

By manipulating GDTs, EDDIE (like other GP algorithms) may handle rule sets of arbitrary size. Besides, rules are easy to understand and evaluate by human users, which makes them more attractive in financial forecasting than neural networks, most of which are black boxes [Goonatilake & Treleaven 1995].

Note that only simple Boolean functions were used in the above example GDT. In principle, there is nothing to stop EDDIE using functions of any level of complexity, though we have only used simple ones. The exact grammar (including the functions) used in our experiments can be found in [Li & Tsang 1999, 2000]. FGP can only create rules by combining the functions provided. The combined rules can be very complex, but not novel.

For a GP to work, one must be able to evaluate each GDT, and assign to it a *fitness* value, which reflects the quality of the GDT. EDDIE maintains a set of GDTs called a *population* and works in iterations. In each iteration, GDTs are picked from the population weighted randomly using *fitness-proportionate selection*, which means that the fitter a GDT is, the greater chance it has of being picked. The set of all GDTs thus picked form a *mating pool* from which pairs of GDTs, which are referred to as *parents*, are picked. A branch in each parent is picked at random. The parents then exchange the subtrees under those branches, as shown in figure 3. This operation is called *crossover*. Offspring are *mutated* occasionally, which is done by replacing random elements of the GDT by random (or heuristically determined) values. The possibly mutated offspring will then replace the old GDTs to form the new population. There are many variations in the way that the population is updated by new offspring, the way that the initial population is generated, the way that parents are picked, the way that crossover and mutation are done, etc. These will not be elaborated here.

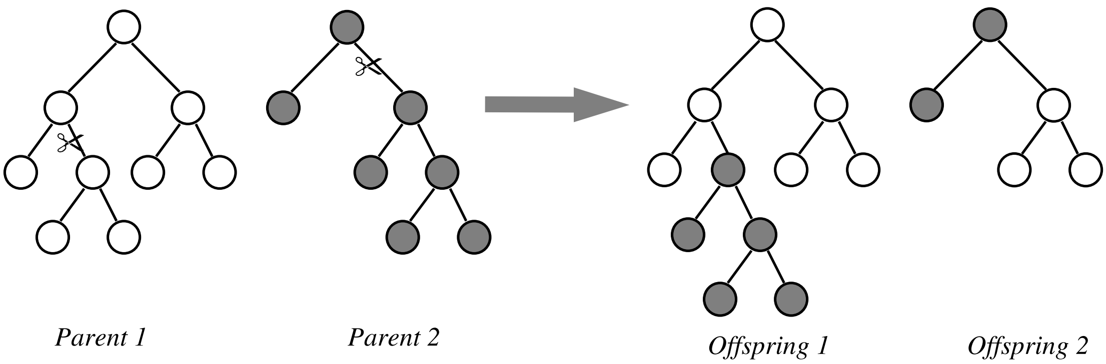

*Figure 3. Crossover in genetic programming (*cut off points marked by ✂)

3. Historical Note: EDDIE-1 in Horse Racing

For completeness, it is worth mentioning the first implementation of EDDIE, called EDDIE-1, was applied to horse racing [Tsang et al 1998]. Trained with 150 records on UK handicap races in 1993, EDDIE-1 beat random decisions (control), handicapper (experts) and the favourite (the publics) in 30 races. The sample size was small (it was only the lack of data that prevented us from doing larger scale tests), and therefore results were interpreted with caution. Nevertheless, it encouraged us to further develop EDDIE for financial forecasting.

It must be emphasized there is no magic in GP. All it can offer is an efficient search engine for the space of possible models, and generate rules that the user can interpret. It is very important for the expert users of EDDIE to define the prediction task and the grammar for describing hypotheses, and the factors to be included in the data set. For example, EDDIE can be asked to make investment recommendations. It can also be asked to find patterns to fit a given set of data [Butler 1997]. More applications will be described in the following sections.

4. EDDIE In Finding Investment Opportunities In Shares and Indices

We used EDDIE to find rules for predicting whether the following goal is achievable at any given day:

Goal *G*: the index will rise by *r%* within *n* trading days.

In FGP-1 (co-named EDDIE-3; FGP stands for *Financial GP*), we used {If-then-else, And, Or, Not, <, =, >} as functions. The crossover operator was designed to take care of the type of branches in the GDTs. Terminals were indicators, thresholds or *conclusions*. Indicators were derived from rules in the finance literature, such as [Alexander 1964; Fama & Blume 1966; Sweeney 1988]. Examples of indicators are:

Filter_63: Today's price over the minimum price of the previous 63 trading days;

TRB_50: Today's price over the maximum price of the previous 50 trading days

TRB_50 is derived from the *Trading Range Breakout* rule [Brock et al 1992]. Thresholds are numbers. Conclusions could be either *Positive* (meaning that G is predicted to be achievable) or *Negative*. We call a trading day a *positive position* if G holds.

FGP-1 and FGP-2 (co-named EDDIE-3, to be elaborated later) have been extensively tested on a wide range of data sets, including S&P 500 (1970-1974, 1995-2000) and Dow Jones Industrial Average (1969-1988), FTSE index (1995-2000), ten individual US shares, about ten UK shares and ten Australian shares up to Year 2000. Some (but not all) of these results were published [Li & Tsang 1999a,b; Tsang et al 1998].

One of the ways in which we evaluated the performance of FGP-1 was to choose *r* and *n* values to make roughly 50% of the test period positive. This allowed us to compare FGP-1 against random decisions without undue complication. Like EDDIE, C4.5 [Quinlan 1993] also generates decision trees. In [Li & Tsang CEC99] FGP-1 produced results statistically better than random decision ("buy" or "not buy" with probability half) and C4.5 in the rate of prediction correctness and precision.

Tests were also performed in bullish, sluggish and bearish markets. We have done enough tests to draw the following conclusions:

1. Patterns have been found in the indices that we tested on. For example, rules trained from 1000 records of the FTSE-100 index (22/8/1994 to 21/5/1998) were used to predict whether the 500 subsequent days (22/5/1998 to 20/4/2000) were positive or negative with r = 4% and n = 22 (days). With appropriate parameter settings, over 65% of the predictions were correct.
2. EDDIE found patterns in some but not all of the shares that we tested on. For example, EDDIE found no patterns in HSBC and BT at London Stock Exchange (1995 to 2000) and many of the Hong Kong and Australian shares (1992 to 1999). This may indicate structural changes in the companies in question, or other regime shifts affecting the macro-economy in the period tested. Clear repeated patterns were found in shares such as BAA, Marks & Spencer and Cadbury at London Stock Exchange between 1995-2000.
3. The question of whether such patterns will repeat themselves in the future will be left to the chartists and the fundamentalists. However, should the patterns found repeat themselves in the future, then they have commercial value.

4. Parameter settings are important. Added to the importance of human input (as explained above), EDDIE should not been seen as an automatic financial decision solution to users who have no knowledge of GP and financial forecasting.
5. The data that we have used so far is very minimal: we have only used daily closing prices in the above-mentioned tests. It is reasonable to believe that better quality predictions can be made with more data input, such as the inclusion of trade volume, interest rates, market trend, etc.

5. EDDIE In Combining Ordinal Predictions

Ordinal data play an important part in financial forecasting. For example, advice from an expert may take the form "buy" or "do not buy". Fan et al [1996] collected the ordinal forecast by nine experts on the weekly movements of the Hong Kong Heng Seng index over 103 weeks. These forecast take the form of “bullish”, “bearish” or “sluggish”. Fan et al demonstrated that better predictions can be made by combining these nine expert opinions. We repeated the same experiment with FGP-1 (co-named EDDIE-3), producing predictions comparable to those by Fan et al in quality [Tsang & Li 1998]. The rules generated by FGP were more accurate than any of the individual expert opinions.

Given that the above database was quite small, we tested FGP-1 on nearly 11 years of daily closing S&P 500 index data (2,700 data cases from 2 April 1963 to 25 January 1974). We asked FGP-1 to make daily buy or not-buy decisions by combining recommendations using technical rules that we managed to find in textbooks [Fama & Blume, 1966; Sweeney, 1988; Brock, *et al.*, 1992]. Similar results were obtained, namely that rules found by FGP-1 was statistically more accurate than the recommendations by any of the input rules. Our experiments demonstrated that FGP-1 could be a useful tool for combining and improving certain ordinal input.

6. EDDIE In Improving Precision

Improving precision is considered important in many applications of machine learning. For example, in financial decision making, the recommendation “not to invest” is often less interesting than the recommendation “to invest”. The former leads to no action. If it is wrong, the user loses an investment opportunity, which may not be serious if other investment opportunities are available. On the other hand, a recommendation to invest leads to commitment of funds. If it is wrong, the user fails to achieve the target rate of return. Precision is defined as the number of real investment opportunities over the total number of predicted opportunities.

One of the objectives in our research is to increase precision, or reduce the rate of failure when EDDIE recommends investing. FGP-2 (co-named EDDIE-4) implements a method for tuning the fitness function to favour degree of precision that reflects the user's preference [Li & Tsang 2000]. This is achieved by introducing a constraint-directed fitness function to FGP.

FGP-2 was tested on historical Dow Jones Industrial Average (DJIA) Index. Data from a seven year period (3,035 trading days, from 07/04/1969 to 11/10/1976) was used for training. Decision trees generated by FGP-2 were tested on data from a three year period (1,900 trading days, from 12/10/1976 to 09/04/1981). For the purpose of analysis, we chose $r = 2.2$ and $n = 21$ days, which gave roughly 50% of positive positions in both the training and test periods. The overall performance of FGP-2 was not impressive in this period. The purpose of our experiments was to show that one could, to a certain extent, tune the precision rate by adjusting a constraint parameter in FGP-2. Higher precision can be achieved at the cost of higher missing opportunities, without affecting the overall accuracy of the system.

To test the robustness of FGP-2, we further partitioned the whole test period into three partitions with downtrend, side-way-trend as well as uptrend periods. We were encouraged by that fact that the rules generated in the same period was accurate in all markets of different trends. Similar results were achieved in forecasting individual shares. Extensive experiments support the fact that FGP-2 allows users to tune the precision and the rate of missed opportunities to suit their needs.

7. EDDIE In Arbitrage

The object of the exercise in Markose et al (2000) is to develop and implement EDDIE on intra daily tick data for stock index options and futures arbitrage in a manner that is suitable for online trading when windows of profitable arbitrage opportunities exist for short periods from one minute to under ten minutes. Following from Tucker (1991), arbitrage based on the Put-Call-Futures (P-C-F, for short) parity condition is used to conduct arbitrage in index options that can bypass the cash leg of the spot stock index market which is prohibitively expensive. As the FTSE-100 spot index has a stock index futures and a

European style index option traded on it, the two derivatives are tailor made for P-C-F arbitrage. Recent work by Markose and Er (2000) on the FTSE-100 stock index options indicates that options that are far from maturity with even fewer than 40 days can present arbitrage opportunities especially for the short P-C-F arbitrage.

Many factors make the application to arbitrage challenging and different from the applications presented in the previous sections. Profitable arbitrage from stock index futures and options typically depend on a number of complex determinants pertaining to the two derivatives and the underlying spot index and also on the speed of detecting profit signals as windows of opportunities are short lived. As the market by its nature does not provide plentiful profitable arbitrage opportunities, the arbitrageur must exploit as many profitable arbitrage opportunities (and avoiding as many loss making positions) as possible. Further, in the implementation of EDDIE to P-C-F arbitrage while considerable analysis went into what constitutes a good GDT for stock index arbitrage, an equal amount of effort had to be put into the pre-processing and embedding of the data so that EDDIE can perform effectively.

In the standard textbook *ex ante* analysis of index arbitrage, the naïve premise is that the arbitrageur waits for a contemporaneous profit signal using the following well-known condition on the violation of a riskless hedge P-C-F no arbitrage profit position. The condition for the short P-C-F arbitrage profit where the arbitrageur sells the overpriced futures at the present discounted value of the futures bid price, F_{bt}, and creates a synthetic long futures position by buying a call option at the call ask, C_{at}, and selling a put option at the put bid, P_{bt}, and lending the present discounted value of the exercise price, $Xe^{-r\tau}$, is given by

$$F_{bt}\, e^{-r\tau} - (C_{at} - P_{bt} + Xe^{-r\tau} + T.C) > 0. \qquad \text{(A)}$$

Here, note that τ is the time to maturity, r is the LIBOR rate (London Interbank Offer Rate) and the T.C denotes the transactions cost which is taken to be 0.1% value of the FTSE-100 index futures contract for the broker/floor market maker. It is held that only this category of traders rather than private investors can conduct the P-C-F arbitrage profitably.

EDDIE is provided with time synchronized price triplets of call, put and futures intra daily tick price data from the LIFFE (January 1991- June 1998) that include both traded and quoted prices. Further, the following time synchronized explanatory variables were used: the exercise price, spot index value, time to Maturity, Type: 0 if the triplet was generated from bid-ask quotations, 1 if it is generated from trade data set and profit after TC as determined in (A). The total sample had 12,513 observations and half of these were used for training EDDIE.

In the textbook *ex ante* analysis the arbitrageur waits for the contemporaneous profit signal (A) and then trades. As profits based on this do not survive more than ten minutes and given execution delays of five minutes, the question we put to EDDIE is:

> *At any point in time corresponding with the occurrence of a matched P-C-F price triplet, how many adjacent 5 minute intervals can it predict as being profitable for arbitrage in a given direction?*

Note the recommended arbitrage positions are judged profitable or not by the criterion given in equation (A).

This application illustrates the importance of human intervention in EDDIE (as described in figure 1). When the original data set, which contained the raw call, put, strike prices, etc., was presented to EDDIE, no reliable patterns were discovered. A number of modifications had to be made to the way in which the explanatory variables were input before EDDIE made reliable recommendations to trade. After six rounds of enhancing the data, theoretically important inputs such as moneyness of the option, the overpricing or underpricing of the futures with respect to the spot index and others were input to EDDIE to improve its performance. The more expert knowledge we channelled into the system (we are now dealing with version 10 of the data set), the better performance we achieved. Current results show that very high precision can be achieved (80% or above), though EDDIE missed 98% of the arbitrage opportunities. As the training data was known to contain some 14% profit opportunities from all observations, our current research focuses on picking up substantially more patterns to cover these profitable arbitrage opportunities. In a forthcoming paper, we shall report EDDIE's trading profitability against the above textbook naïve strategy and a number of expert rules developed in Markose and Er (2000).

Acknowledgement

Full credit should be given to James Butler who invented the acronym EDDIE. He implemented EDDIE-1 and tested it on horse racing. Many colleagues and students at University of Essex have contributed to the EDDIE project through discussion. In particular, we would like to thank Paul Scott, Jeff Reynolds, Robert Pug and Tony Lawson (MSc student). This project is partly supported by two Research Promotion Funds by the University of Essex. Jin Li was supported by the Overseas Research Scholarship and the University of Essex Scholarship. We have benefited from the insightful comments by the anonymous referee.

References

1. Alexander, S.S., *Price movement in speculative markets: Trend or random walks, No. 2*, in Cootner, P. (ed.), the random character of stock market prices, MIT Press, Cambridge, MA, 1964, 338-372
2. Angeline, P. & Kinnear, K.E.Jr. (ed.), *Advances in genetic programming II*, MIT Press, 1996
3. Brock, W., Lakonishok, J. & LeBaron, B., *Simple technical trading rules and the stochastic properties of stock returns*, Journal of Finance, 47, 1992, 1731-1764
4. Butler, J.M., *EDDIE beats the market, data mining and decision support through genetic programming*, Developments, Reuters Limited, Vol.1, July 1997
5. Chen, S-H. & Yeh, C-H., *Speculative trades and financial regulations: simulations based on genetic programming*, Proceedings of the IEEE/IAFE 1997 Computational Intelligence for Financial Engineering (CIFEr), New York City, March 1997, 123-129
6. Chen, S-H. & Yeh, C-H, *Toward a Computable Approach to the Efficient Market Hypothesis: An Application of Genetic Programming*, Journal of Economic Dynamics and Control, Vol.21, 1997, pp.1043-1063
7. Fama, E.F. & Blume, M.E., *Filter rules and stock-market trading*, Journal of Business 39(1), 1966, 226-241
8. Fan, D.K., Lau, K-N. & Leung, P-L. *Combining ordinal forecasting with an application in a financial market*. Journal of Forecasting, Vol. 15, No.1, Wiley, January 1996, 37-48
9. Goonatilake, S. & Treleaven, P. (ed.), *Intelligent systems for finance and business*, Wiley, New York, 1995

10. Kaboudan, M.A., Evaluation of forecasts produced by genetically evolved models, Proceedings, 6th International Conference on Computing in Economics and Finance, Society for Computational Economics, Barcelona, July 2000
11. Koza, J.R. *Genetic Programming: on the programming of computers by means of natural selection*. MIT Press, 1992
12. Koza, J., Goldberg, D., Fogel, D. & Riolo, R. (ed.), *Procedings, First Annual Conference on Genetic programming*, MIT Press, 1996
13. Li, J. & Tsang, E.P.K, (1999a), Improving technical analysis predictions: an application of genetic programming, *Proceedings of The 12th International Florida AI Research Society Conference*, Orlando, Florida, May 1-5, 1999, 108-112.
14. Li, J. & Tsang, E.P.K, Investment decision making using FGP: a case study, *Proceedings of Congress on Evolutionary Computation (CEC'99)*, Washington DC, USA, July 6-9 1999.
15. Li, J. & Tsang, E.P.K, *Reducing failure in investment recommendations using genetic programming*, Computing in Economics and Finance Conference, Barcelona, July 2000 (revised version submitted to Journal of Computational Economics)
16. Malkiel, B., *Efficient market hypothesis*, in Newman, P. , Milgate, M. & Eatwell, J. (eds.), New palgrave dictionary of money and finance, Macmillan, London 1992
17. Markose, S., Tsang, E.P.K, Er, H. & Salhi, A., *Put-call-futures index arbitrage trading with evolutionary decision trees*, Manuscript, Institute of Studies in Finance, University of Essex, 2000.
18. Markose, S. & Er, H., *The black effect and cross market arbitrage in FTSE-100 Index futures and options*, Manuscript, Department of Economics, University of Essex, 2000.
19. Neely, C., Weller, P. & Ditmar, R., Is technical analysis in the foreign exchange market profitable? a genetic programming approach, in Dunis, C. & Rustem, B. (ed.), Proceedings, Forecasting Financial Markets: Advances for Exchange Rates, Interest Rates and Asset Management, London, May 1997
20. Quinlan, J.R., *C4.5: programs for machine learning*, Morgan Kaufmann, San Mateo, 1993
21. Sweeney, R.J., *Some new filter rule test: Methods and results*, Journal of Financial and Quantitative Analysis, 23, 1988, 285-300
22. Tsang, E.P.K., Li, J. & Butler, J.M., EDDIE beats the bookies, *International Journal of Software, Practice & Experience*, Wiley, Vol.28 (10), 1998, 1033-1043
23. Tsang, E.P.K. & Li, J., Combining Ordinal Financial Forecasts with Genetic Programming, Manuscript, submitted to Journal of Forecasting, 1998 (under revision)
24. Tucker, A.L., *Financial Futures, Options and Swaps*, West Publishing Company, St. Paul, MN., 1991