Grammatical Evolution in Finance and Economics: A Survey

Anthony Brabazon

Natural Computing Research and Applications Group,
School of Business, University College Dublin, Ireland.
anthony.brabazon@ucd.ie

Summary. Finance was one of the earliest application domains for Grammatical Evolution (GE). Since the first such study in 2001, well in excess of 100 studies have been published employing GE for a diverse range of purposes encompassing financial trading, credit-risk modelling, supply chain management, detection of tax non-compliance, and corporate strategy modelling. This chapter surveys a sample of this work and in doing so, suggests some future directions for the application of GE in finance and economics.

1 Introduction

The application of computational algorithms whose design is metaphorically derived from phenomena in the natural world to finance and economics has providence dating back some thirty years. A huge literature running to thousands of papers has resulted. Previous survey articles spanning these contributions include [24, 11]. There have also been several edited volumes covering this area [22, 23, 25, 26, 34, 35, 36]. A good introduction to the application of biologically inspired algorithms to various aspects of financial modelling is provided in [19].

The late 1980s and early 1990s saw a plethora of studies applying (initially) neural networks for financial forecasting purposes and (slightly later) evolutionary approaches, particularly genetic algorithms (GAs). In the latter case, attention was primarily focused on the application of GAs for model parameter optimisation and variable selection [8, 44]. As evolutionary automatic programming (EAP) methodologies such as Genetic Programming (GP) were introduced [58, 60] they too were applied by practitioners and researchers for economic modelling, financial forecasting [56], and trading system induction [4]. It is an interesting footnote in the history of GP that one of the earliest exemplar applications of the methodology by John Koza was to recover the well-known *exchange equation* $M = PQ/V$ which relates the money supply (M), price level (P), gross national product (Q), and velocity of money (V)

in an economy [58, 59]. Somewhat quaintly, the paper notes that a then state-of-the-art Apple Macintosh II PC was used to access the training data for the study!

The first finance-related papers which employed a Grammatical Evolution (GE) methodology were published in 2001 arising quite shortly after GE's introduction in 1998 [76]. With the benefit of hindsight it is perhaps unsurprising that finance was one of the early application domains for GE given the stream of related research which was arising in GP around the same time. In common with GP, much of the early GE research in finance was proof of concept in nature, being constrained by the availability of computing power and data (relative to the abundance of both resources today). Nonetheless, this work laid the foundation for more comprehensive subsequent studies which also took advantage of the increasing maturity of GP and GE methodologies. By the current day, application of GP and GE approaches in finance have moved beyond academic studies to include implementation in multiple instances by industry practitioners.

Structure of Chapter

In this chapter a survey of the literature which has applied GE to finance and economics since 2001 is provided, along with some suggestions for future research. As with any survey, there are difficult choices in deciding which research contributions to include and which to omit. By design, a broad coverage of the relevant literature is provided and this precludes detailed discussion of individual studies.

The rest of this chapter is organised as follows. In the following sections we introduce the application of GE across a range of financial applications including financial trading (Sect. 2), credit-risk modelling (Sect. 3), supply chain management (Sect. 4.1), detection of tax non-compliance (Sect. 4.2), and corporate strategy modelling (Sect. 4.3). Section 5 concludes this chapter, outlining some opportunities for future work.

2 Financial Trading

The earliest financial applications of GE were for the purposes of financial trading [28, 67, 68]. A variety of trading approaches are seen in practice and we outline three of these below in order to place the relevant GE literature into context.

Fundamental Analysis

Taking the example of investing in a share, investment using a fundamental analysis approach concentrates on the use of accounting and other information about a company, as well as industry and macro-economic data, in order to identify companies which are mispriced by the market. In other words, the object is to identify shares which are good value (underpriced by the market), or shares which are overpriced by the market (and therefore are candidates for 'shorting').

To apply this approach, the investor needs to develop stock screening rules in order to decide which shares to select. The utility of these rules can be tested using historical data, with the best rule (or set of rules) then being used for investment purposes. EAP methods such as GP or GE can be applied to evolve the actual structure of the filter rules. In spite of the widespread use of fundamental analysis (including the application of EAP methods) by actively-managed investment funds, few academic studies have explored the application of EAP for the task of filter rule development.

Technical Analysis

Under a technical analysis approach, investors attempt to identify imbalances in the supply and demand for a financial asset using information from the time-series of the asset's trading price and volume. Technical indicators (pre-processed price and volume time series data about a financial asset) can be used in isolation, or combined, in order to produce a 'trading signal'. For example, one technical indicator that technical analysts could consider is the *moving average convergence-divergence (MACD) oscillator*, calculated by taking the difference of a short-run and a long-run moving average. If the difference is positive, it is taken as a signal that the market is trending upward, with a buy signal being generated when the shorter moving average crosses the longer moving average in an upward direction. A sell signal is generated in a reverse case. Therefore, a sample MACD trading rule is:

$$\textit{IF } x\text{-day MA of price} \geq y\text{-day MA of price}$$

$$\textit{THEN Go Long ELSE Go Short}$$

where $x < y$ (for example $x = 10$ and $y = 50$). The MACD oscillator is a crude band-pass filter, removing both high-frequency price movements and certain low-frequency price movements, depending on the precise moving average lags selected.

A trader who wishes to construct a trading system using technical indicators as inputs faces several decisions, namely:

1. Which indicators will be used?
2. What parameter values (lag periods used to calculate the indicator and threshold values for the indicator to trigger a trading action) should be used?
3. How should the indicators be combined to produce a trading signal?

This results in a high-dimensional search space which is suitable for EAP methodologies.

Arbitrage

Arbitrage approaches to trading aim to profit by exploiting price differences between financial instruments, trading on different markets or trading in different forms. For example, suppose a share traded on one stock exchange for $23.78 and on another for $23.82, an investor could arbitrage by buying at the lower price and selling simultaneously at the higher price. As would be expected, such a simple arbitrage opportunity would tend to be closed very quickly and transactions costs of buying and selling can negate apparent arbitrage possibilities.

There are a plethora of arbitrage trading strategies employed in financial markets. One exemplar, based on *Put-Call Parity*, is illustrated in [84]. The concept underlying this trade is that the price of a 'long' position on an asset and an associated put (the right to sell that asset in the future at a specified price) must be equal to the price of a long call on the same asset and a long position in a risk-free bond. If either the put or call option are mispriced the investor can, in theory, make a risk-free gain by constructing a portfolio of the four financial instruments. The above example, describes an arbitrage opportunity between the cash market (for the asset) and the options market. Arbitrage opportunities can also exist between cash and futures markets and between futures and options markets.

GE and Trading

From the above discussion it is evident that there are a variety of trading approaches to which model induction approaches could be applied. Thus far, the bulk of applications of GE, and indeed other model induction methodologies, have adopted a technical analysis perspective. Within this perspective, the trading time horizon can be varied from inter-day to high-frequency trading, with choice of trading horizon impacting on the temporal resolution of input data required by the trading system. A multiplicity of factors come into play in operationalising trading systems in the real-world including market microstructure, market structure, money management, the risk appetite of the investor and their trading 'style'. Academic contributions rarely consider all of these factors in detail.

2.1 Applications of GE for Trading System Design

The earliest study applying GE for trading system design [67] focussed on the UK FTSE 100 index from the period 26/4/1984 to 4/12/1997. Initially, attention was restricted to the generation of a moving average inter-day trading system. The evolved system produced a buy, sell or 'do nothing' signal with trades being left open for a pre-defined, fixed period of 10 days. The fitness function was defined as trading profit less maximum drawdown in order to encourage the generation of trading rules with monotonic equity curves. Allowance was made in the calculation of fitness for both trading cost and slippage. Despite the simplicity of the trading system specification, it outperformed a naive buy-and-hold benchmark whilst maintaining a smaller 'at risk' average investment in the market. Follow on studies extended the approach to encompass an extended range of technical indicators [28] and an extended range of markets (DAX, ISEQ, Madrid stock exchange and NIKKEI) [68, 38, 66].

In [45] a number of methodological improvements to the canonical trading system of [67] were implemented. Instead of implementing a hard 'all or nothing' investment threshold, the trading system was allowed to invest a variable amount depending on the strength of the generated trading signal. A variety of population replacement strategies were also implemented in order to assess the impact of varying these on system performance. Changing this parameter in the GE implementation altered the importance of memory in the evolutionary process by impacting on the speed of convergence of the population. The resulting trading systems generated quite stable performance between in/out of sample periods and outperformed the original system when tested on US market data. After taking account of trade costs the system did not outperform a buy-and-hold benchmark however.

When applying an evolutionary approach to generate trading rules it is important to embed as much useful domain knowledge as possible in order to bias the generation process towards syntactically correct, well-formed, trading rules. For example, it would not be sensible to directly compare technical indicators of differing types (consider a moving average value for price which could assume a large or small value and a stochastic indicator which by definition is limited to values between 0 and 100) in constructing a trading rule. Certain indicators are also only validly parameterised within defined ranges. These issues require careful grammar design and an exemplar of this is provided in the GE implementation of [31].

A key aspect of financial markets is that they are dynamic. Consequently, trading systems need to be adaptive to changing conditions. In the early applications of GE for financial trading, a fixed in sample training, out of sample testing, approach was adopted. An obvious limitation of this approach is that trading rules will have a 'shelf life', with their performance deteriorating over time during the out of sample period. This problem was addressed in [46], implementing a 'live trading' system which updated its trading rules in real-time as new information became available. An initial training period is set aside on which the population of proto-trading rules is trained, with the aim that a competent population is evolved after a certain number of generations, G. The system then goes 'live'. The trading system takes the best performing rule from the initial training period, and uses this rule to trade for each of the following x days. After x days have elapsed, the training window moves forward in the time series by x days, and the population is retrained over the new data window for a number of generations g, where $g < G$. This embeds both a memory and an adaptive potential in the trading system, as knowledge of good past trading rules is not completely lost, rather it serves as a starting point for their subsequent adaptation. The size of g is crucial. A small value of g means that memory is emphasised over adaptation, as the new data has relatively less chance to influence the trading rules. This could be considered a tuning parameter that could be used to alter the adaptive characteristics of the system. In [47] the performance of an adaptive trading system based on the above approach was compared with that of a system which is periodically retrained ab initio, with the adaptive system demonstrating better results.

The bulk of trading systems using an EAP methodology have employed a technical analysis approach to trading. One notable exception to this is provided by [39, 40] which combine technical, fundamental and macro-economic analysis in a hybrid, top-down, paradigm. Initially, prospective investments are screened using macro-economic and corporate data, with the final invest decision in selected companies being made based on technical analysis. GE is used to create both the stock selection and trading rules respectively, thereby implementing a dual-layer version of GE. The system is found to produce better results with lower drawdown than benchmark systems including a GE approach which solely relies on technical indicators as inputs.

Intraday Data

High-frequency financial data corresponds to data that is sampled at small time intervals (i..e, at high-frequency) during the trading day. Therefore, high-frequency financial data, in the limit, can contain every single event on an order book during a trading day. Traditionally, high-frequency data was difficult, and costly, to obtain so most early research examining the application of various computational intelligence techniques for trading used end of day data (i.e., a single data point for each trading day).

The first application of GE to high-frequency financial data was that of [12] which applied a technical analysis approach to price and volume data for Ford and IBM, sampled on a five-minute frame. Each data vector included the opening and closing prices, the high and low price, and the volume traded, for that five-minute interval. All open positions were closed out before the end of each day, with no inventory of stock being carried overnight.

In a complete implementation of a technical analysis trading system, the value of technical indicators are used to determine when to enter and exit a trade (the exit strategy can also incorporate a stop loss trigger in order to manage downside risk). Earlier implementations of technical analysis trading systems usually adopted a simpler design, such as having a fixed exit horizon with no stop loss trigger, and therefore only in effect evolving an entry point.

In addition to evolving the entry strategy for investment, [12] also assessed the performance of three different exit strategies namely, *standard close*, *ex-*

tended close and *stop-loss, take-profit close*. In the standard close, the evolved systems automatically close out trading positions 30 minutes after they are opened. In the extended close, the system rechecks after 30 minutes whether the prediction is unchanged from the initial prediction, and if it is, the trade is extended for a further 30 minutes. In the stop-loss, take-profit close, the position is initially held for 30 minutes, and thereafter, if the position generates a loss of 0.1% it is closed immediately, and profit is automatically taken on any position which makes a profit of 0.8% by closing the position once the take-profit trigger is hit.

In both the training and test periods, the extended close exit strategy was found to outperform the standard close strategy and the stop-loss, take-profit strategy, with all three strategies significantly outperforming a buy and hold benchmark, albeit that trading costs were ignored in this study. This highlights the impact that the choice of exit strategy can have on the returns produced by a trading system.

A paper by [2], building on earlier related work by the same authors [1], takes a more comprehensive approach, using GE to coevolve technical rules for entry/exit (for both long and short positions) and also implements a stop loss mechanism, using data based on five-minute bars, for a stock listed on the London stock exchange. An earlier paper by [77] had employed a similarly sophisticated approach in evolving trading strategies using GP on high-frequency tick data of the USD/EUR exchange rate.

Other studies to evolve both entry and exit rules include that of [51] which applied GE to develop a trading system for the S&P 500 E-mini index futures contract. This future is traded on the CME. The system used a simple moving average and a relative strength indicator. A moving window approach was applied to train, validate and test the trading strategies in order to create a method suitable for an online trading system. The authors noted the ease with which domain knowledge could be incorporated into the grammar in order to evolve targeted trading strategies, resulting in relatively transparent trading rules which were capable of human interpretation.

One notable feature of the literature applying computational intelligence approaches (including EAP methodologies) to develop trading systems is an increasing integration between the fields of finance and computer science. In particular, some recent contributions have exhibited a stronger input from the domain of finance, drawing much more deeply on finance theory. A good exemplar of this work is provided by [64] which develops a high frequency trading system using GE and applies this to NASDAQ limit order book data (rebuild order book data). The trading system is motivated by a theoretical liquidity asymmetry theorem from the market microstructure literature, with GE being used to exploit volume inefficiencies at the bid-ask spread. GE evolves a condition which describes volume distributions in the order book. If the condition is violated, the system places a limit order on the side of the market where the volume is too small, expecting the temporary price impact to be reversed. In this study, attention is focussed on long-strategies only (i.e., the trading system aims to exploit liquidity undersupply on the bid side of the market). The resulting system is found to be able to generate profitable and robust strategies.

Foreign Exchange Markets

The use of technical analysis for trading spot foreign exchange markets is also common so an obvious extension of the early work applying GE for trading purposes was to apply it to these markets. The initial pilot study [21] applied the canonical approach of [28] to daily Daily US $ - STG exchange rates for the period 9/3/93 to 13/10/97. This work was extended to include the US $ - DM and US $ - Yen exchange pairs with results indicating that the developed rules earned positive returns in hold-out sample test periods after allowing for trading and slippage costs [14, 18]. Subsequent studies focussing on foreign-exchange markets include [77], already discussed above.

Sentiment Analysis

Traditionally, stock trading models only incorporated quantitative data, drawn from the market, financial statements or macro-economic data. One source of information which has attracted increasing attention as an input into trading models in recent years is text data drawn from either internet message boards, social media or the financial press.

While initial studies such as [82] looked at raw message count information, the next step was to consider the content or 'sentiment' of these messages in order to assess whether investors are (un)favorably disposed towards a stock. Sentiment analysis has become a major area of research covering natural language processing, computational linguistics and text mining.

Using text data in quantitative models has been made easier by the commercial availability of 'tagged' databases of financial news. One example of these is the Dow Jones Elementised News Feed which places discrete pieces of news - keywords, timestamps, symbols and other crucial data - into XML-tagged fields for easy parsing and direct embedding into trading programs.

A study by [61] used a post-processed version of this data (where the stories were classed as positive, negative or neutral, in essence a metric of market sentiment, with respect to a particular market), combined with a GP methodology, to predict intraday price jumps on the S&P 500 up to an hour before they occur. The results indicated that the system was successfully able to predict stock price movement using the news stories alone, without access to market price or volume data. Future work using text data will doubtless seek to extract more detailed semantic meaning from news stories as inputs to trading systems.

Fitness Function Design

One of the critical decisions in applying GE, or indeed any EAP approach, to trading system generation is the choice of fitness function. Poor choices will produce poor trading systems. For example, selection of raw profit as a fitness metric is quite likely to lead to the generation of trading systems with undesirable risk characteristics, in terms of the variance or skewness of trading returns.

A trading strategy will generate a corresponding distribution of payoffs (returns) if implemented repetitively, with the shape of the distribution depending on the nature of the strategy. The first moment of the underlying probability density function corresponds to the expected payoff to the strategy per trade with the higher moments of the distribution (variance, skewness, and kurtosis etc.) describing how payoffs vary around this value. In other words, the higher moments assess the risk that returns may vary from expected.

In most implementations of EAP to trading system design, attention is focussed on the first and second moments of the distribution (i.e., the mean and variance of returns) with the aim being to maximise the former and minimise the latter. However, there are clear limitations of this approach. It implicitly assumes constant risk aversion as it treats positive and negative outcomes symmetrically. Of course, the design of the fitness function can be much more nuanced than merely maximising a metric such as the *Stirling ratio*:

$$\frac{\text{Return}}{\text{Drawdown}} \tag{1}$$

and can be specified so as to bias the generated trading system towards systems with a desired distribution of returns.

Another problem with traditional approaches to fitness function selection is that risk metrics such as variance do not take account of the temporal ordering of returns. A sequence of negative returns can lead to large drawdowns which can have detrimental consequences if an individual investor or fund runs out of capital, suggesting an important role for consideration of all aspects of the shape of the produced equity curve in assessing trading systems.

A study examining the behavior of GE generated trading models evolved using different choices of fitness functions by [30] illustrated that these choices have a very significant impact. This study forms a useful starting point for future work concerning fitness function design.

One issue which impacts on all trading system induction methods is that of data snooping. In essence, when a dataset is used multiple times for model selection an apparently good model could occur due to chance alone rather than representing a truly robust model of the data-generating process [87]. In these cases, the trading system will likely perform poorly out of sample. Data snooping is a particular concern in powerful methodologies such as EAP approaches due to the very large number of trading systems that can be generated and tested against the same dataset during training.

The study of [6] addressed this issue, aiming to investigate the profitability of evolved trading rules, having controlled for data-mining bias. The approach taken implements a multi-criterion fitness function that in addition to a measure of profitability, takes into account Hansen's Superior Predictive Ability test, which can directly quantify the effect of data-mining bias, by testing the performance of the best mined rule in the context of the full universe of technical trading rules. Another study addressing this issue, in the context of a for-ex trading system, is [78] which constructs a framework for trading rule selection using a-priori robustness strategies, where robustness is gauged on the basis of time-series bootstrap and multi-objective criteria.

2.2 Trade Execution

As noted above, most academic contributions on trading systems omit detailed consideration of implementation issues. In real-world trading, especially high-frequency trading, careful attention must be paid to the structure of the market being traded and the different ways in which investors can interact with this market.

Market Structure

Most large financial markets now operate as an electronic double auction limit order book on which investors can post buy or sell orders at a specific, desired, price. These orders are known as limit orders and are visible to market participants. Limit orders can be cancelled at any time prior to their execution. Alternatively, an investor may submit a market order which is executed immediately at current prices. The choice between use of a limit or a market order depends on the sensitivity of the investor to the probability of order execution versus the price paid.

As limit orders are visible when placed on the order book, investors wishing to buy or sell a large quantity of stock will usually seek to manage the trade in order to minimise its price impact. If an investor places a large 'buy' order, potential sellers will mark up prices. An obvious strategy in response is to attempt to break up the order into smaller pieces and push it out to the market a bit at a time. The potential drawback of this is that the market may move against the buyer, hence design of an execution strategy tries to balance market impact versus the risk of not filling the order. An efficient trade execution strategy seeks to address the following:

- Timing – when should the order be placed and / or what interval of time should there be between orders (what is the schedule?)
- Type – should the order be a market, limit, reserve, hidden order?
- Sizing – what size order should be sent to the market?
- Pricing – at what price should the order be, aggressive or passive?

- Destination – there are many market destinations and types. Which one will provide the best conditions of execution for the order?
- Management – if a limit order has been submitted, how should this order be managed post submission?

A variety of fitness functions could be designed to drive the evolution of the execution strategy but a common metric of trade execution performance is its *Volume Weighted Average Price* (VWAP):

$$VWAP = \frac{\sum(Price \cdot Volume)}{\sum(Volume)}$$

The VWAP of a strategy can be calculated and benchmarked against (for example) the overall VWAP for that share during the period of the trading strategy's execution. The aim is to evolve a strategy which produces as competitive a VWAP as possible.

GE and Trade Execution

A novel approach was taken by [41, 42] where GE was used to evolve a dynamic trade execution strategy, with the resulting rule adapting to changing market conditions. Based on the finance literature analysing the relationship between order placement and the information content of limit order books, six order book metrics were selected as potential inputs for an execution strategy.

A practical issue arises in the assessment of evolved execution strategies, in that they cannot be easily backtested using historical data as it is very difficult to assess the impact that an execution strategy would have produced ex ante. Apart from this issue, another practical problem is that historical order book information only represents a single sample path through time and hence, using this information to estimate the likely future utility of any specific execution strategy is problematic.

In [41, 42] the training and evaluation of all trade execution strategies was implemented in an artificial limit order market, simulated using an agent-based model, parameterised using data drawn from real-world financial markets. By implementing an artificial stock market environment, it is possible to create a closed world which allows the testing of new execution strategies over multiple runs, potentially allowing us to develop robust execution strategies.

In the implementation, GE was found to be able to evolve quality trade execution strategies and its results proved highly competitive against two basic benchmark execution strategies. A detailed discussion of the application and the relevant background finance literature is provided in [43].

There is notable scope for further research utilising GE for trade execution. One obvious route is to widen the number of market variables which can be included in the evolved execution strategies.

The above are not the only studies which have applied GE in an artificial stock market setting. In [86], a Maslov limit-order model which can be parameterised to generate controlled cyclic behaviour in the price signal is implemented. A trader, whose strategy is evolved via GE using a range of technical indicators, interacts with this market and may place limit or market orders. The objective of the study is to gain insight into the evolved trader behaviour and discover how this alters with changes in the cyclic behaviour of the market. Potentially such understanding could provide insights into how best to generate trading strategies which need to generalise to a market with a range of cyclic behaviours. In essence, the study uses the artificial stock market as a tunable model in order to provide a closed environment for the testing of trading strategies.

3 Credit Risk Modelling

The assessment of credit risk plays an important role in lending decisions. While the precise nature of credit facilities can vary depending on the agreement between the borrower and lender, in all cases lenders need to assess the capability of a borrower to make both interest and capital repayments over the lifetime of the loan. The decision as to whether to extend a loan and if so extended, its pricing, will depend on this assessment of credit risk.

In the case of consumers, a credit risk model could consider factors such as current income, age, occupation, current employment status, past borrowing record and so on [85, 88]. Corporate risk models could include factors such as data drawn from the financial statements of the firm, data drawn from financial markets (such as share price), general macro-economic data, and non-financial firm-specific information. The development of these risk assessment models requires the discovery of suitable explanatory variables and model form, with model output being a metric of credit / default risk.

Bankruptcy Prediction

A closely related research topic is that of bankruptcy prediction where the objective is to predict whether a firm will declare bankruptcy within a predetermined forecast horizon. This is typically styled as a classification problem, the object being to correctly predict the classification of a firm out of sample, as being solvent or bankrupt. The pioneering study in this domain was that of Altman (1968) [5] in which five accounting ratios were selected and then combined to produce a linear discriminant classification model for corporate bankruptcy. A Z score was calculated for each company and this value determined whether the company was classified as likely to go bankrupt or likely to remain solvent. The original Altman classifier had the form:

$$Z=0.012X_1+0.014X_2+0.033X_3+0.006X_4+0.999X_5$$

where:

X_1 = working capital to total assets

X_2 = retained earnings to total assets
X_3 = earnings before interest and taxes to total assets
X_4 = market value of equity to book value of total debt
X_5 = sales to total assets

Subsequently, more sophisticated statistical approaches including logit and probit regression [65, 92] were utilised. As the range of computational intelligence techniques for classification have expanded, each new technique has been applied in turn to credit scoring and corporate failure prediction with model induction methodologies such as artificial neural networks and support vector machines producing good results. A drawback of these approaches is that the resulting classifers are generally not human-readable which can preclude their use in practice as some jurisdictions require lenders to justify decisions not to grant loans. Hence, EAP methodologies can be useful.

The earliest application of GE for the purposes of corporate failure prediction was [69] which explored the potential of GE to uncover rules to assist in predicting corporate failure using information drawn from financial statements of 178 publicly quoted US firms, drawn from the period 1991 to 2000. Twenty two financial ratios, drawn from prior finance literature on corporate failure prediction, were supplied as potential explanatory variables. In the initial study, the grammar was restricted to generate linear classifiers. The results obtained were competitive against other classification methods with good classification accuracies being obtained out of sample. An interesting aspect of the evolved classifiers, which were separately evolved for each prediction horizon, was that they indicated a clear 'trajectory towards failure', with low profits and high interest payments as a percentage of profits being particular risk factors two and three years prior to failure, with short-term liquidity issues arising as a key risk factor in the final year before a firm's demise.

Years Prior to Failure	In Sample	Out Of Sample
1	85.9%	80%
2	82.8%	80%
3	75.8%	70%

Table 1. The accuracies reported for each of the three years prior to failure based on best evolved GE classifier.

A key assumption in the above study (and in all other literature on corporate failure prediction) is that the selected financial ratios provide the optimal pre-processing of raw numbers from the financial statements of corporates. The GE methodology provides an easy way to address this issue as raw data drawn from the financial statements can be provided as inputs instead of pre-processed financial ratios, with the grammar allowing GE to create classifiers consisting of 'self-evolved' ratios. This approach was adopted in [13]. In essence therefore, the human domain knowledge that was supplied in [69] in the form of pre-selected financial ratios, was omitted in this study.

The out of sample predictive accuracies obtained were similar to those in [69] indicating that not alone could GE generate quality linear classifiers for the problem at hand but it could also recover the domain knowledge embedded in the finance literature concerning financial ratios with good information content for prediction of bankruptcy. An expanded version of this study was subsequently presented in [15]. Other relevant work is that of [3] which considers the issues of unbalanced datasets, a particular feature of bankruptcy prediction modelling, as the number of failing companies is typically relatively small in comparison with the number of solvent companies.

While the above studies did not make use of non-financial information about the firms, or general macro-economic data, it would not be difficult to extend the approach to include such information.

Bond Rating Prediction

When large corporates wish to raise debt which is tradable on a recognised financial market, they need to obtain a bond rating from an independent rating firm such as Standard & Poor's (S&P) or Moody's. The bond rating firms undertake an assessment of either the proposed lender's general credit-worthiness (an issuer credit rating) or an assessment of the the credit-worthiness of a specific bond issue they propose to make (a bond issue credit rating). Therefore the ratings serve as a surrogate measure of the risk of non-payment of interest or capital. These ratings impact on the borrowing cost and the marketability of issued bonds.

The ratings are revised periodically as the circumstances of the borrower change. Being able to anticipate bond rating changes could potentially provide a useful input into a stock or bond trading model. It would also provide useful information for pricing of credit-risk derivatives concerning that borrower.

Following some initial work in the 1960s, there was increased research interest in attempting to predict corporate bond ratings from the 1980s [49, 50, 52]. In common with corporate failure prediction a feature of bond rating prediction is that there is no unambiguous theoretical framework for guiding the choice of explanatory variables, or model form. Rating agencies assert that their credit rating process involves consideration of both financial and non-financial information about the firm and its industry, but the precise factors utilised, and the related weighting of these factors, are not publicly disclosed by the agencies. In the absence of an underlying theory, most published work on bond rating prediction employs a data-inductive modelling approach, using firm-specific financial data as explanatory variables, in an attempt to recover the model used by the rating agencies. This produces a high-dimensional combinatorial problem.

The initial application of GE to bond rating prediction focussed on issuer-credit ratings and aimed to predict whether a rating would fall into investment-grade or junk category [16, 17, 20]. Financial data, and the associated Standard & Poor's issuer-credit ratings of 600 public US firms, drawn from the years 1999/2000 were used to train and test the model.

The best developed model was found to be able to discriminate in-sample (out-of-sample) between investment-grade and junk bond ratings with an average accuracy of 87.59 (84.92)% across a five-fold cross validation, producing very similar results to an MLP applied to the same data. In contrast to the MLP models, the GE classification model was reasonably compact and produced a human-readable classification rule which concorded with domain knowledge. Given that GE was restricted to evolve linear classification rules, the comparability of its results with those from application of an MLP would indicate that the relationships between financial data and resulting bond ratings are not, in fact, highly non-linear.

Other Related Problems in Finance

There are several other areas of research in finance which seek to predict a corporate outcome using financial and other information about a firm and its industry. A significant stream of work involves the prediction of targets for merger or takeover. Obviously, being able to accurately predict which firms may be a target in advance of the market generally, could provide useful trading information. A wide a variety of methodologies have been applied in an attempt to forecast takeover and merger targets, including univariate analysis [74], MDA [7, 80], probit / logit analysis [62, 73], multi-layer perceptrons (MLPs)[32] and self-organising maps [55].

Another related stream of work is the prediction of an auditor's 'going-concern' qualification of the financial statements of a company. As this qualification states that the auditor does not believe that the company will continue in existence, the issuing of such a qualification will typically have a major impact on the decisions of investors, bank lenders, creditors and employees. When an entity receives a going concern qualification they will usually suffer serious repercussions including restrictions on trade credit, constraints on the raising of further finance/capital, and the possibility of a share price collapse. This domain has attracted research attention over the past three decades. A wide variety of methodologies have been applied in an attempt to predict going concern qualification including univariate analysis, LDA [63], logit [9], probit [48] and MLPs [83]. Most studies have relied heavily on the use of company accounting data as modelling inputs. As yet, GE has not been applied to either of these areas.

4 Other Finance and Economics Applications of GE

Over the years, GE has been introduced to a variety of application areas in finance and economics outside of the areas already discussed in this chapter. In this section we briefly overview a few of these.

4.1 Supply Chain Management

Supply chain management concerns the management of the flow of goods and services from point of origin to point of consumption. It encompasses flows of raw materials, work-in-process, and finished goods. Efficient and effective management of these flows is crucial in order to ensure the requisite availability of products / services to each customer as needed, and in order to control the costs of the supply chain as a whole. Modern supply chains can be very complex ranging over multiple countries and encompassing many different organisations.

A particular challenge in managing supply chains is that unexpected events can occur such as unanticipated changes in final customer demand, unexpected events in production systems (for example, closure of a factory), or changes in legislation impacting on the flow of some product. Such unexpected events can result in a 'bullwhip effect' with the initial impact being magnified at other stages in the supply chain, producing inefficiencies and high costs. In the worst-case scenario, customer service declines, lead-times increase, sales are lost, costs go up and capacity is adjusted [71]. A common example of the bullwhip effect is where orders to a supplier in the supply chain have a larger variance than sales to the buyer in the same supply chain, referred to as demand distortion. This can occur when suppliers do not have good information on final sales and over-respond to small fluctuations in these.

A traditional problem in supply chain management is to derive the optimal ordering policy for an individual firm, given the information available to it. Many approaches have been taken in the literature to derivation of this policy with [71] novelly applying GE in an attempt to derive an optimal ordering policy for agents in a multi-tier supply chain. In this study, GE is implemented in a simulation environment where artificial agents are playing the Beer Game [79]. In this game, shipments arrive from upstream players, orders arrive from downstream players, orders are filled and shipped where possible, affecting the inventory and backorders of a player, the player in each step of the game (i.e., each time period) decides how much to order to replenish their inventory, and finally inventory holding costs and backorder costs are calculated for each player every week [71]. Different scenarios can be generated, depending on whether the final demand from customers is deterministic or stochastic. The GE grammar is defined so as to allow the generation of a valid ordering policy and the costs of a specific policy can be estimated by running a simulation of final demands and actions of each agent in the supply chain. The results indicated that GE was capable of producing effective ordering policies. Developments on this initial study are discussed in [70] and [72].

4.2 Tax Non-Compliance Detection

The setting of tax evasion can be considered as a co-evolutionary arms race in which the tax evader is attempting to outwit the rules as enforced by the tax authority. The latter is attempting to catch evaders via tax audits and to close off loopholes in tax regulation. A novel application of GE was suggested by [53] in the simulation modelling of this setting. In the study, a search heuristic called STEALTHis implemented that can simulate the co-evolution of abusive tax avoidance schemes and audit scores, where GE is applied to generate a series of transactions for the purposes of this simulation. The system has potential use in guiding policy formulation as it allows the exploration of the likely forms of (new) tax schemes in response to changes in audit policies of tax authorities. Obviously, such information by anticipating abusive tax avoidance schemes could potentially be used to help select between different forms of audit procedures in order to counteract these. The developed system is quite complex and full details are provided in [54] and [75]. On a broader level, this work provides another exemplar of the application of GE in an agent-based setting.

4.3 Corporate Strategy

An enduring research question is whether organisational strategy matters in terms of explaining corporate performance, and if so, how much does it matter. While it appears self-evident at first glance that corporate decisions regarding firm scope and resource allocation, would affect firm performance the empirical evidence is not so clear cut, with studies of the performance of diversification and retrenchment strategies producing varying results [10, 81]. The strategy domain is characterised by a lack of strong theoretical frameworks. It is also notable that unlike areas such as financial market prediction, there have been virtually no applications of EAP methodologies to questions in the strategy domain.

An initial study by [29] introduced GE to this domain, modelling the relationship between corporate strategy and shareholder wealth. A shareholder perspective is adopted in the study and it is assumed that the success of corporate strategy decisions is judged by equity markets. To allow for exogenous factors which could impact on stock market values generally, a relative performance metric, a market-value-added (MVA) rank is utilised (MVA is defined as: Market Value of Firm - Original Value of Capital Invested) and the performance of the firm is determined by whether it improved its MVA ranking in the Stern-Stewart 1000 listing (published in Fortune magazine) over a four year period.

A total of 430 US firms were selected from the Stern-Stewart Performance 1000 list for the study and sixteen potential explanatory variables, which can proxy strategic intent [81], drawn from their financial statements are collected from the Compustat database for each firm. The best classifier correctly categorised the direction of performance ranking change in 66% of the firms in the training set and 65% in the out-of-sample validation set providing support for a hypothesis that changes in corporate strategy are linked to changes in corporate performance. A detailed discussion of the results and the form of the evolved classifiers can be found in [29]. Scope exists to extend the application of GE into this domain beyond this initial pilot study.

5 The Future

GE has found a variety of applications in finance and economics since its introduction in 1998. It is notable that the sophistication of the studies employing the methodology has increased over the years moving from initial proof of concept studies to much more robust, industry-strength, applications. This process has been facilitated by the dissemination of knowledge about GE beyond computer science to include domain experts from finance and economics. A key selling point for many finance academics and practitioners concerning GE is its ability to encapsulate domain knowledge in the grammar. This helps reassure users that the resulting models are plausible. Hence, the areas with greatest potential for the application of GE are those where we have good data concerning the phenomenon of interest but only a partial understanding of how this data might fit together.

Financial trading will continue to be an important area of application for GE given its natural fit with this application. We can expect to see increasing sophistication in these application with fuller implementation of smart entry and exit strategies, and greater attention being paid to market structure. There is also substantial opportunity to undertake work concerning the design of appropriate fitness functions for trading applications. We can expect to see greater integration of non-financial information, such as that from social media or official news wires into trading systems.

The range of financial instruments traded on markets has expanded enormously over the past twenty years, moving far beyond the trading of shares and debt instruments to encompass a wide variety of financial derivatives. A key issue for investors wishing to trade in derivatives for speculative or hedging purposes is the determination of a fair price for the derivative. While GP has been applied for the purposes of reverse-engineering pricing models ([33, 37, 57, 89, 91] is a sampling of this work) and to develop hedging strategies [90], the powerful flexibility of domain knowledge incorporation into a grammar in GE makes these obvious areas for future work.

Another area with significant potential for the application of GE is that of agent-based modelling (ABM). Although there have already been a number of studies applying GE in an ABM framework ([41, 42, 43, 53, 54, 86] being a sampling of these), it is perhaps surprising that GE has not gained greater traction as a tool for ABM. The nature of GE makes it particularly amenable for application in ABM as it is relatively easy for modellers to place a desired structure via the grammar definition on the strategies that agents can employ, while still allowing considerable room for agents to adapt their strategies. In addition to modelling of agents in financial markets, there are a multiplicity of opportunities for policy-focussed research via application of a GE methodology to ABM in economics.

The application of GE to finance and economics is nearly 'fiche bliain ag fás'.[1] It will be fascinating to see the continued development of this work over the next twenty years.

References

1. Adamu, K., Phelps, S. (2009) Modelling financial time series using grammatical evolution. In: Proceedings of the Workshop on Advances in Machine Learning for Computational Finance, London, UK
2. Adamu, K., Phelps, S. (2011) Coevolution of technical trading rules for high frequency trading. In: Electrical Engineering and Applied Computing, Lecture Notes in Electrical Engineering, Vol. 90, pp. 311-322, Springer, Berlin
3. Alfaro-Cid, E., Cuesta-Canada, A., Sharman, K. and Esparcia-Alcazar, A. (2008). Stong typing, variable reduction and bloat control for solving the bankruptcy prediction problem using genetic programming. In: Natural Computing in Computational Finance, pp. 161-186, Springer, Berlin.
4. Allen, F. and Karjalainen, R. (1999). Using genetic algorithms to find technical trading rules. Journal of Financial Economics, 51:245-271.
5. Altman, E. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. Journal of Finance, 23:589-609.
6. Agapitos, A., ONeill, M. and Brabazon, A. (2010). Evolutionary Learning of Technical Trading Rules without Data-Mining Bias, International Conference on Parallel Problem Solving from Nature (PPSN 2010), LNCS 6238, pp. 294-303, Springer, Berlin.
7. Barnes, P. (1990). The prediction of takeover targets in the U.K. by means of multiple discriminant analysis, Journal of Business Finance & Accounting, 17(1):73-84.
8. Bauer R. (1994). Genetic Algorithms and Investment Strategies, Wiley, New York.
9. Bell, T.B. and Tabor, R.H. (1991). Empirical Analysis of Audit Uncertainty Qualifications, Journal of Accounting Research, 29(2):350-370.
10. Bowman, E. and Helfat, C. (2001). Does Corporate Strategy Matter?, Strategic Management Journal, 22:1-23.
11. Brabazon, A., Dang, J., Dempsey, I. and O'Neill, M. (2011). Natural Computing in Finance: A Review, in Handbook of Natural Computing: Theory, Experiments and Applications G. Rozenberg, T. Baeck and J. Kok (eds), pp. 1707-1735, Berlin: Springer.

[1] Fiche bliain ag fás, or translated from Irish into English - 'Twenty Years A-Growing', is the title of a famous autobiographical book written by Muiris Ó Súilleabháin in the Irish language. The book is set in the Great Blasket Island which lies off the south west coast of Ireland, part of a group of islands inhabited until 1953 by a completely Irish-speaking population. The book forms part of an 'island literature' which details the end of a Gaelic way of life in the early 20th century.

12. Brabazon, A., Meagher, K., Carty, E. , O'Neill, M. and Keenan, P. (2005a). Grammar-mediated time-series prediction, Journal of Intelligent Systems, 14(2-3):123-143.
13. Brabazon, A. and O'Neill, M. (2003). Anticipating Bankruptcy Reorganisation from Raw Financial Data using Grammatical Evolution, Proceedings of EvoIASP 2003: Applications of Evolutionary Computing, LNCS 2611:, pp. 368-378, Springer, Berlin.
14. Brabazon, A. and O'Neill, M. (2003a). A Grammar model for foreign-exchange trading, Proceedings of the International Conference on Artificial Intelligence 2003 (ICAI '03), pp. 492-499: CSEA Press.
15. Brabazon, A. and O'Neill. M. (2004a). Diagnosing Corporate Stability using Grammatical Evolution, International Journal of Applied Mathematics and Computer Science, 14(3):363-374.
16. Brabazon, A. and O'Neill, M. (2004b). Bond-Issuer Credit Rating with Grammatical Evolution, Proceedings of EvoIASP 2004: Applications of Evolutionary Computing, LNCS 3005:, edited by Raidl, G., Cagnoni, S. et al., pp. 268-277, Springer, Berlin
17. Brabazon, A. and O'Neill, M. (2004c). A Grammatical Approach to Bond-Issuer Credit Rating, Proceedings of OR 2004, Tilburg 1-3 September 2004.
18. Brabazon, A. and O'Neill, M. (2004d). Evolving technical trading rules for foreign-exchange markets, Computational Management Science, 1(3-4):311-327.
19. Brabazon, A. and O'Neill, M. (2006). Biologically Inspired Algorithms for Financial Modelling, Springer, Berlin.
20. Brabazon, A. and O'Neill, M. (2006b). Credit Classification Using Grammatical Evolution, Informatica, 30(3):325-335.
21. Brabazon, A., O'Neill, M. and Ryan, C. (2002a). Trading Foreign Exchange Markets Using Evolutionary Automatic Programing, In: Proceedings of the First Workshop on Grammatical Evolution: GECCO 2002, pp. 133-137, New York, July 9 - 13, 2002.
22. Brabazon, A. and O'Neill, M. (eds.) (2008). Natural Computing in Computational Finance, Springer, Berlin.
23. Brabazon, A. and O'Neill, M. (eds.) (2009). Natural Computing in Computational Finance (Volume II), Springer, Berlin.
24. Brabazon, A., O'Neill, M. and Dempsey, I. (2008). An Introduction to Evolutionary Computation in Finance, IEEE Computational Intelligence Magazine, 3(4):42-55.
25. Brabazon, A., O'Neill, M. and Maringer, D (ed.). (2010) Natural Computing in Computational Finance (Volume III), Springer, Berlin.
26. Brabazon, A., O'Neill, M., Maringer, D (eds.). (2011) Natural Computation in Computational Finance (Volume IV), Springer, Berlin.
27. Brabazon A, O'Neill M, McGarraghy S (2015) Natural Computing Algorithms, Springer, Berlin.
28. Brabazon, A., O'Neill, M., Ryan, C. and J. Collins (2001). Developing a Market Timing System Using Grammatical Evolution, Proceedings of AAANZ 2001, Auckland, July 1-3, 2001.
29. Brabazon, A., O'Neill, M., Ryan, C. and Matthews, R. (2002) Evolving Classifiers to Model the Relationship Between Strategy and Corporate Performance Using Grammatical Evolution, Proceedings of EuroGP 2002, In: Pro-

ceedings of EuroGP 2002, Genetic Programming, LNCS 2278:, pp. 103-113, Springer, Berlin.
30. Bradley, R., Brabazon, A. and O'Neill, M. (2010). Objective Function Design in a Grammatical Evolutionary Trading System, Proceedings of the 2010 IEEE Congress On Evolutionary Computation, pp. 3487-3494, IEEE Press.
31. Bradley, R., Brabazon, A. and O'Neill, M. (2010a). Evolving Trading-rule Based Policies, Proceedings of EvoFin 2010, Applications of Evolutionary Computation, LNCS 6025, pp. 250-259, Springer, Berlin.
32. Cheh, J.J., Weinberg, A., Yook, K.C., (1999) An Application Of an Artificial Neural Network Investment System To Predict Takeover Targets, The Journal of Applied Business Research, 15(4):33-44.
33. Chen, S-H., Lee, W-C. and Yeh, C-H. (1999). Hedging derivative securities with genetic programming. International Journal of Intelligent Systems in Accounting, Finance and Management, 8(4):237-251.
34. Chen, S-H. (eds.) (2002). Genetic Algorithms and Genetic Programming in Computational Finance, Kluwer Academic Publishers.
35. Chen, S-H. (eds.) (2002). Evolutionary Computation in Economics and Finance, Physica-Verlag.
36. Chen, S-H., Wang, P. and Kuo, T-W. (eds) (2007). Computational Intelligence in Economics and Finance (vol. II), Springer, Berlin.
37. Chidambaran, N., Lee, C. and Trigueros, J. (1998). Adapting Black-Scholes to a non-Black-Scholes environment via genetic programming. Proceedings of the IEEE/IAFE/INFORMS 1998 Conference on Computational Intelligence for Financial Engineering (CIFEr), pp. 197-211, IEEE Press.
38. Contreras I., Hidalgo J.I., Núñez-Letamendia L. (2013) Combining Technical Analysis and Grammatical Evolution in a Trading System. In: Applications of Evolutionary Computation. EvoApplications 2013. Lecture Notes in Computer Science, vol 7835, pp. 244-253, Springer, Berlin
39. Contreras, I., Hidalgo J.I. and Núñez-Letamendia, L. (2017). A hybrid automated trading system based on multi-objective grammatical evolution, Journal of Intelligent & Fuzzy Systems, 32(3):2461–2475.
40. Contreras, I., Hidalgo J.I., Núñez-Letamendia, L. and J. Manuel Velasco, J. (2017). A meta-grammatical evolutionary process for portfolio selection and trading, Genetic Programming and Evolvable Machines, DOI:10.1007/s10710-017-9304-1
41. Cui, W., Brabazon, A. and O'Neill, M. (2010). Evolving Dynamic Trade Execution Strategies Using Grammatical Evolution, Proceedings of EvoFin 2010, Applications of Evolutionary Computation, LNCS 6025, pp. 191-200, Springer, Berlin.
42. Cui, W., Brabazon, A. and O'Neill, M. (2010a). Evolving Efficient Limit Order Strategy using Grammatical Evolution, Proceedings of the 2010 IEEE Congress On Evolutionary Computation, pp. 2408-2413, IEEE Press.
43. Cui, W., Brabazon, A. and O'Neill, M. (2011). Adaptive Trade Execution using a Grammatical Evolution Approach, International Journal of Financial Markets and Derivatives, 2(1/2):4-31.
44. Deboeck, G. (1994). Trading on the Edge: Neural, Genetic, and Fuzzy Systems for Chaotic Financial Markets, John Wiley and Sons, New York.
45. Dempsey, I., O'Neill, M. and Brabazon, A. (2002). Investigations into Market Index Trading Models using Evolutionary Automatic Programming, Proceedings of AICS 2002, LNAI 2464, pp. 165 - 170, Springer, Berlin.

46. Dempsey, I. O'Neill, M. and Brabazon, A. (2004). Live Trading with Grammatical Evolution, Proceedings of the Third Workshop on Grammatical Evolution: GECCO 2004, Seattle, June 26 - 30, 2004.
47. Dempsey, I. O'Neill, M. and Brabazon, A. (2006). Adaptive Trading with Grammatical Evolution, Proceedings of the Congress on Evolutionary Computation (CEC 2006), Vancouver, July 16-21, 2006, pp. 9137-9172, IEEE Press: New Jersey.
48. Dopuch, N., Holthausen, R.W., and Leftwich R.W. (1987). Predicting Audit Qualifications with Financial and Market Variables, The Accounting Review, LXII(3):431-454.
49. Dutta, S. and Shekhar, S. (1988). Bond rating: a non-conservative application of neural networks. Proceedings of IEEE International Conference on Neural Networks, II, pp. 443-450, IEEE Press.
50. Ederington, H. (1985). Classification models and bond ratings. Financial Review, 20(4):237-262.
51. Gabrielsson, P., Johansson, U. and König, R. (2014). Co-evolving online high-frequency trading strategies using grammatical evolution, In: Proceedings of the 2014 IEEE Conference on Computational Intelligence for Financial Engineering and Economics (CIFEr 2014), pp.1-8, IEEE Press.
52. Gentry, J., Whitford, D. and Newbold, P. (1988). Predicting industrial bond ratings with a probit model and funds flow components. Financial Review, 23(3):269-286.
53. Hemberg, E., Rosen, J., Warner, G. and O'Reilly U-M (2015) Tax Non-Compliance Detection Using Co-Evolution of Tax Evasion Risk and Audit Likelihood, In: Proceedings of the 15th International Conference on AI and Law (ICAIL 2015), pp. 79–88, ACM: New York.
54. Hemberg, E., Rosen, J., Warner, G. et al. (2016). Detecting tax evasion: a co-evolutionary approach, Artificial Intelligence and Law, 24(2):149–182.
55. Hickey, R, Little, E. and Brabazon, A. (2006). Identifying Merger and Takeover Targets Using a Self-Organising Map, Proceedings of the 2006 International Conference on Artificial Intelligence (ICAI 06), Las Vegas 26-29 June 2006, Vol. 1, pp. 408-413, CSEA Press.
56. Iba H. and Nikolaev N. (2000). Genetic programming polynomial models of financial data series. Proceedings of CEC 2000, pp. 1459-1466, IEEE Press.
57. Keber, C. (2000). Option Valuation with the Genetic Programming Approach. Proceedings of the sixth international conference, pp. 689-703, Cambridge, MA: MIT Press.
58. Koza, John R. (1990) Genetic Programming: A Paradigm for Genetically Breeding Populations of Computer Programs to Solve Problems. Stanford University Computer Science Department Technical Report STAN-CS-90-1314. June 1990.
59. Koza, John R. (1990b) A genetic approach to econometric modeling. Paper presented at Sixth World Congress of the Econometric Society, Barcelona, Spain. August 27, 1990.
60. Koza, J. (1992). Genetic Programming, Massachusetts: MIT Press.
61. Larkin, F. and Ryan C. (2008). Good news: using news feeds with genetic programming to predict stock prices. In: O'Neill, M. et al (eds.), Proceedings of the 11^{th} European Conference on Genetic Programming (EuroGP 2008), pp. 49-60, LNCS 4971, Springer.

62. Meador, A., Church, P. and Rayburn, L. (1996). Development of Prediction Models for Horizontal and Vertical Mergers, Journal of Financial And Strategic Decisions, 9(1):11-23.
63. Mutchler, J.F. (1985). A Multivariate Analysis of the Auditor's Going-Concern Opinion Decision, Journal of Accounting Research, 23(2):668-682.
64. Oesch, C. and Maringer, D. (2016). Low-latency liquidity inefficiency strategies, Quantitative Finance, 17(5):717–727.
65. Ohlson, J. (1980). Financial ratios and the probabilistic prediction of bankruptcy. Journal of Accounting Research, 18(1):109-131.
66. O'Neill, M., Brabazon, A., Ryan, C. (2002). Forecasting Market Indices using Evolutionary Automatic Programming: A case study, in Genetic Algorithms and Genetic Programming in Economics and Finance, Shu-Heng Chen (ed.), pp. 174-195, Kluwer Academic Publishers.
67. O'Neill, M., Brabazon, A., Ryan, C. and Collins J. (2001). Evolving Market Index Trading Rules Using Grammatical Evolution, Proceedings of EvoIASP 2001, in LNCS: Applications of Evolutionary Computing, pp. 343-353, Springer, Berlin,.
68. O'Neill, M., Brabazon, A., Ryan, C. and J. Collins (2001a). Developing a Market Timing System using Grammatical Evolution, in Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2001), San Francisco, USA, July 7-11, 2001, pp. 1375-1381, Morgan Kaufmann.
69. O'Neill, M., Brabazon, A., Ryan, C. and R. Matthews (2002). Grammatical Evolution and Corporate Failure Prediction, in Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2002), New York, USA, July 9-13, 2002, pp. 1011-1019, Morgan Kaufmann.
70. Phelan, M. (2016). Bullwhips & beer: grammatical evolution in supply chains! SIGEVOlution, 9(4):3–8, ACM: New York
71. Phelan, M. and McGarraghy, S. (2007). Mitigating the Bullwhip Effect in Supply Chains using Grammatical Evolution, In: System Dynamics Society eds. System Dynamics Society Conference Boston, Massachusetts, USA, `https://www.systemdynamics.org/conferences/2007/proceed/papers/PHELA331.pdf`
72. Phelan, M. and McGarraghy, S. (2016). Grammatical evolution in developing optimal inventory policies for serial and distribution supply chains. International Journal of Production Research, 54(1):336–364.
73. Powell, R. (1995). Takeover Prediction in the UK: A Multilogit Approach, Working Papers in Accounting and Finance, Accounting and Finance Division, The Queen's University of Belfast, 95(3).
74. Rege, U., (1984). Accounting Ratios to Locate Take-Over Targets, Journal of Business Finance and Accounting, 11(3):301-311.
75. Rosen, J., Hemberg, E. and OReilly U-M (2016). Dynamics of Adversarial Co-evolution in Tax Non-Compliance Detection, In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2016), pp. 1087–1094, ACM: New York.
76. Ryan, C., Collins, J.J., O'Neill, M. (1998). Grammatical Evolution: Evolving Programs for an Arbitrary Language. Proceedings of the First European Workshop on GP, pp. 83-95, Springer, Berlin.
77. Saks, P., Maringer, D. (2009) Evolutionary Money Management. Natural Computing in Computational Finance, Volume 293 of the series Studies in Computational Intelligence, pp. 169-190, Springer, Berlin

78. Schmidbauer, H., Rösch, A., Sezer, T. et al. (2014) Robust trading rule selection and forecasting accuracy, Journal of Systems Science and Complexity, 27:169–180.
79. Sterman, J. (1989). Modeling Managerial Behavior - Misperceptions of Feedback in a Dynamic Decision-Making Experiment, Management Science, 35(3):321–339.
80. Stevens, D. (1973). Financial Characteristics of Merged Firms: A Multivariate Analysis, Journal of Financial and Quantitative Analysis, 8(2):149-159.
81. St. John, C., Balakrishnan, N. and Fiet, O. (2000). Modelling the relationship between corporate strategy and wealth creation using neural networks, Computers and operations research, 27:1077-1092.
82. Thomas, J. and Sycara, K. (2002). GP and the Predictive Power of Internet Message Traffic, in Genetic Algorithms and Genetic Programming in Computational Finance, Shu-Heng Chen (ed)., pp. 81-102, Kluwer Academic Publishers.
83. Thompson, D., Thompson, S. and Brabazon, A. (2007). Predicting Going Concern Audit Qualification using Neural Networks, Proceedings of the 2007 International Conference on Artificial Intelligence (ICAI 07), Las Vegas 25-28 June 2007, Vol. 1, pp. 199-204, CSEA Press.
84. Tung, W. and Quek, C. (2005). GenSoOPATS: a brain-inspired dynamically evolving option pricing model and arbitrage system. Proceedings of the IEEE International Conference on Evolutionary Computation (CEC 2005), pp. 1722-1729, IEEE Press.
85. West, D. (2000). Neural network credit scoring models. Computers and Operations Research, 27:1131-1152.
86. Whigham, P. and Withanawasam, R. (2011). Evolving a robust trader in a cyclic double auction market, In: Proceedings of the 13th annual Conference on Genetic and Evolutionary Computation (GECCO 2011), pp. 1451–1458, ACM: New York.
87. White H (2000). A reality check for data snooping, Econometrica, 68:1097-1126.
88. Yobas, M, Crook, J. and Ross, P. (2000). Credit scoring using neural and evolutionary techniques. IMA Journal of Mathematics Applied in Business and Industry, 11:111-125.
89. Yin, Z., Brabazon, A. and O'Sullivan, C. (2007). Adaptive genetic programming for option pricing. Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2007), pp. 2588-2594, ACM Press.
90. Yin, Z., Brabazon, A., O'Sullivan, C., O'Neill, M. (2015). A Genetic Programming Approach for Delta Hedging IEEE Congress on Evolutionary Computing Sendai, Japan, pp. 3312-3318, IEEE Press.
91. Yin, Z., O'Sullivan, C. and Brabazon, A. (2016). An Analysis of the Performance of Genetic Programming for Realised Volatility Forecasting. Journal of Artificial Intelligence and Soft Computing Research, 6(3):155-172.
92. Zmijewski, M. (1984). Methodological Issues Related to the Estimation of Financial Distress Prediction Models, Journal of Accounting Research - Supplement, pp. 59-82.