


Quantum Computing

Neural Networks

Evolutionary Computing

DNA Computing



Anthony Brabazon · Michael O'Neill

# Biologically Inspired Algorithms for Financial Modelling

 Springer

NATURAL COMPUTING SERIES


# Natural Computing Series

Series Editors: G. Rozenberg

Th. Bäck A.E. Eiben J.N. Kok H.P. Spaink

Leiden Center for Natural Computing

---



Advisory Board: S. Amari G. Brassard K.A. De Jong

C.C.A.M. Gielen T. Head L. Kari L. Landweber T. Martinez

Z. Michalewicz M.C. Mozer E. Oja G. Păun J. Reif H. Rubin

A. Salomaa M. Schoenauer H.-P. Schwefel C. Torras

D. Whitley E. Winfree J.M. Zurada

Anthony Brabazon · Michael O'Neill

# Biologically Inspired Algorithms for Financial Modelling

With 92 Figures and 39 Tables

 Springer

*Authors*

Dr. Anthony Brabazon  
University College Dublin  
Belfield, Dublin 4  
Ireland  
anthony.brabazon@ucd.ie

Dr. Michael O'Neill  
University College Dublin  
Belfield, Dublin 4  
Ireland

*Series Editors*

G. Rozenberg (Managing Editor)  
rozenber@liacs.nl

Th. Bäck, J.N. Kok, H.P. Spaink  
Leiden Institute of Advanced  
Computer Science  
Leiden University  
Niels Bohrweg 1  
2333 CA Leiden, The Netherlands

A.E. Eiben  
Vrije Universiteit Amsterdam  
The Netherlands

Library of Congress Control Number: 2005936099

ACM Computing Classification (1998): F.1, F.2, I.2.1, I.2.8, I.6, J.1, J.4

ISBN-10 3-540-26252-0 Springer Berlin Heidelberg New York

ISBN-13 978-3-540-26252-7 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable for prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media  
springer.com

© Springer-Verlag Berlin Heidelberg 2006  
Printed in Germany

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Cover Design: KünkelLopka, Werbeagentur, Heidelberg

Typesetting: by the Authors

Production: LE-TeX Jelonek, Schmidt & Vöckler GbR, Leipzig

Printed on acid-free paper

45/3142/YL - 5 4 3 2 1 0

*To Maria  
Tony*

*To Gráinne  
Michael*

---

## Preface

The field of biologically inspired computing has advanced rapidly over the past decade. One offshoot of this progress has been the development of a large family of biologically inspired algorithms. Broadly speaking, biologically inspired algorithms draw metaphorical inspiration from diverse sources, including the operation of biological neurons, processes of evolution, models of social interaction amongst organisms, and natural immune systems, in order to develop tools for solving real-world problems. This book provides an introduction to a broad range of biologically inspired algorithms and illustrates how they can be applied for financial modelling using a series of case studies. These cases include the modelling of financial markets, the development of financial trading systems, the creation of solvency prediction systems, and the creation of credit rating models. In this book particular emphasis is placed on evolutionary methodologies, particularly the novel, powerful, modelling methodology Grammatical Evolution. No prior knowledge of either biologically inspired algorithms or financial modelling is assumed.

We hope that this book will help spark new ideas in the minds of readers to encourage them to undertake their own work in financial modelling using biologically inspired methodologies.

October 2005

*Anthony Brabazon*  
*Michael O'Neill*

---

## Acknowledgment

We would like to thank several people who contributed to the writing of this book. Thomas Randles and David Brennan helped collect the corporate failure and bond-rating datasets used in some of the case studies. We also thank John McCallig, Sean McGarraghy and David Edelman for useful discussions on several aspects of the material in the text which helped clarify our thoughts (all errors are of course our responsibility!). In particular we would like to thank a number of people who co-authored some of the case study chapters with us: Ian Dempsey who co-authored the adaptive trading case study (Chapter 14); Peter Keenan, Katrina Meagher and Edward Carty who were co-authors of the intra-day case study (Chapter 15); Yue Xi and Qiang Han who were co-authors of the corporate failure ant-model case study (Chapter 18); and Peter Keenan, Alice Delahunty and Denis O'Callaghan who contributed to the AIS bond-rating chapter (Chapter 20). We also extend our thanks to Ronan Nugent of Springer-Verlag for his encouragement of this project, and for his advice on early drafts of the manuscript.

*Anthony Brabazon  
Michael O'Neill*

---

# Contents

<b>1</b>	<b>Introduction</b> .....	1
1.1	Biologically Inspired Algorithms .....	2
1.1.1	Artificial Neural Networks .....	2
1.1.2	Evolutionary Computation .....	2
1.1.3	Social Systems .....	4
1.1.4	Artificial Immune Systems .....	4
1.2	Computer Trading on Financial Markets .....	4
1.3	Challenges in the Modelling of Financial Markets .....	5
1.3.1	Do Prices Follow a Random Walk? .....	6
1.3.2	Attack of the Anomalies .....	7
1.4	Linear Models .....	8
1.5	Structure of the Book .....	10

---

## Part I Methodologies

---

<b>2</b>	<b>Neural Network Methodologies</b> .....	15
2.1	A Taxonomy of NNs .....	15
2.2	The Multi Layer Perceptron .....	16
2.2.1	Training an MLP .....	20
2.2.2	Practical Issues in Training MLPs .....	23
2.2.3	Recurrent Networks .....	28
2.3	Radial Basis Function Networks .....	29
2.4	Self-organising Maps .....	32
2.4.1	Implementing a SOM .....	33
2.5	Summary .....	35
<b>3</b>	<b>Evolutionary Methodologies</b> .....	37
3.1	Genetic Algorithm .....	37
3.1.1	Canonical GA .....	40
3.1.2	Example of the GA .....	41



3.1.3	Extending the Canonical GA	43
3.1.4	Schema and Building Blocks	48
3.2	Differential Evolution	49
3.2.1	DE Algorithm	49
3.3	Genetic Programming	54
3.3.1	More Complex GP Architectures	58
3.4	Combining EA and MLP Methodologies	63
3.5	Applying EAs to Evolve Trading Rules	68
3.6	Recent Developments in Evolutionary Computation	70
3.7	Summary	71
<b>4</b>	<b>Grammatical Evolution</b>	<b>73</b>
4.1	Grammatical Evolution	73
4.1.1	Biological Analogy	74
4.1.2	Mapping Process	76
4.1.3	Mapping Example	79
4.2	Mutation and Crossover in GE	82
4.3	Recent Developments in GE	84
4.3.1	Search Engine	84
4.3.2	Meta-grammars	85
4.3.3	$\pi$ GE	87
4.3.4	Applications and Alternative Grammars	87
4.4	Summary	88
<b>5</b>	<b>The Particle Swarm Model</b>	<b>89</b>
5.1	PSO Algorithm	89
5.1.1	Constriction Coefficient Version of PSO	92
5.1.2	Parameter Settings for PSO	93
5.2	Discrete PSO	94
5.3	Comparing PSO and the GA	94
5.4	MLP-Swarm Hybrids	95
5.5	Grammatical Swarm	95
5.6	Example of a Financial Application of PSO	96
5.7	Recent Developments in PSO	96
5.8	Summary	97
<b>6</b>	<b>Ant Colony Models</b>	<b>99</b>
6.1	Ant-Foraging Models	99
6.1.1	Ant-Foraging Algorithm	100
6.2	A Financial Application of ACO	104
6.3	Ant-Inspired Classification Algorithms	105
6.4	Hybrid Ant Models	105
6.5	Summary	106

<b>7</b>	<b>Artificial Immune Systems</b> . . . . .	107
7.1	Overview of Natural Immune Systems . . . . .	108
7.1.1	Innate vs Adaptive Immunity . . . . .	108
7.1.2	Components of the Immune System . . . . .	108
7.2	Designing Artificial Immune Algorithms . . . . .	113
7.2.1	Negative Selection Algorithm . . . . .	113
7.2.2	Clonal Expansion and Selection Algorithm . . . . .	114
7.3	Financial Application of the Negative Selection Algorithm . . . . .	116
7.4	Summary . . . . .	118

---

## Part II Model Development

---

<b>8</b>	<b>Model Development Process</b> . . . . .	121
8.1	Project Goals . . . . .	121
8.1.1	What to Forecast? . . . . .	121
8.1.2	What Performance Measure Is Appropriate? . . . . .	123
8.2	Data Collection . . . . .	124
8.2.1	Trading Philosophy . . . . .	124
8.2.2	How Much Data Is Enough? . . . . .	128
8.3	Selecting and Preprocessing the Data . . . . .	130
8.3.1	Selection . . . . .	130
8.3.2	Preprocessing . . . . .	130
8.4	Postprocessing the Output . . . . .	134
8.4.1	Entry Strategy . . . . .	134
8.4.2	Exit Strategy . . . . .	134
8.4.3	Money Management . . . . .	135
8.5	Validating the System . . . . .	135
8.6	Implementation and Maintenance . . . . .	140
8.7	Summary . . . . .	142
<b>9</b>	<b>Technical Analysis</b> . . . . .	143
9.1	Technical Indicators . . . . .	144
9.1.1	Moving Average . . . . .	146
9.1.2	Momentum . . . . .	148
9.1.3	Breakout . . . . .	149
9.1.4	Stochastic Oscillators . . . . .	150
9.1.5	Volume Data . . . . .	152
9.1.6	Other Indicators . . . . .	153
9.2	Using Technical Indicators in a Trading System . . . . .	154
9.3	Summary . . . . .	155

---

**Part III Case Studies**

---

<b>10 Overview of Case Studies</b> .....	159
<b>11 Index Prediction Using MLPs</b> .....	161
11.1 Methodology .....	162
11.1.1 Model Selection .....	166
11.1.2 Model Stacking .....	167
11.2 Results .....	169
11.2.1 RMSE and Correlation .....	169
11.2.2 Trading System .....	171
11.3 Discussion .....	172
<b>12 Index Prediction Using a MLP-GA Hybrid</b> .....	175
12.1 Methodology .....	175
12.1.1 Model Construction .....	176
12.2 Results .....	178
12.2.1 MLP-GA .....	179
12.2.2 Analysis of Weight Vectors .....	180
12.3 Discussion .....	182
<b>13 Index Trading Using Grammatical Evolution</b> .....	183
13.1 Methodology .....	183
13.1.1 GE System Setup .....	188
13.2 Results .....	189
13.3 Discussion .....	190
<b>14 Adaptive Trading Using Grammatical Evolution</b> .....	193
14.1 Introduction .....	193
14.2 Methodology .....	193
14.2.1 Moving Window .....	194
14.2.2 Variable Position Trading .....	194
14.2.3 Return Calculation .....	195
14.3 Results .....	196
14.3.1 Training Returns .....	197
14.3.2 Out-of-Sample Returns .....	199
14.4 Discussion .....	201
<b>15 Intra-day Trading Using Grammatical Evolution</b> .....	203
15.1 Background .....	203
15.2 Methodology .....	204
15.2.1 Trading System .....	206
15.2.2 GE System Setup .....	207
15.3 Results .....	208
15.4 Discussion .....	210

<b>16 Automatic Generation of Foreign Exchange Trading Rules</b>	211
16.1 Background	211
16.2 Methodology	212
16.3 Results	214
16.3.1 US-STG	216
16.3.2 US-Yen	217
16.3.3 US-DM	217
16.4 Discussion	218
<b>17 Corporate Failure Prediction Using Grammatical Evolution</b>	219
17.1 Background	220
17.1.1 Definition of Corporate Failure	220
17.1.2 Explanatory Variables	221
17.2 Methodology	222
17.2.1 GE System Setup	223
17.2.2 LDA Method	224
17.3 Results	224
17.3.1 Form of the Evolved Classifiers	225
17.4 Discussion	226
<b>18 Corporate Failure Prediction Using an Ant Model</b>	229
18.1 Background	229
18.2 Methodology	230
18.2.1 Ant System	231
18.3 Results	235
18.4 Discussion	238
<b>19 Bond Rating Using Grammatical Evolution</b>	239
19.1 Background	240
19.1.1 Rating Process	240
19.2 Methodology	241
19.3 Results	243
19.4 Discussion	247
<b>20 Bond Rating Using AIS</b>	249
20.1 Methodology	249
20.2 Results	252
20.3 Discussion	252
<b>21 Wrap-up</b>	255
<b>References</b>	257
<b>Index</b>	271

## Introduction

Over the last decade, a considerable literature on biologically inspired algorithms (BIA) has emerged. These powerful algorithms can be used for prediction and classification, and have clear application for use in financial modelling and in the development of trading systems. Financial markets represent a complex, ever-changing, environment in which a population of investors are competing for profit. Biological entities have long inhabited such environments, and have competed for resources to ensure their survival. It is natural to turn to algorithms which are inspired by biological processes to tackle the task of survival in a financial jungle.

The primary objectives of this book are twofold: to provide readers with an up-to-date introduction to a broad range of BIAs, and to illustrate by means of a series of case studies how the algorithms can be applied for the purposes of modelling financial markets, for the prediction of corporate failure, and for the prediction of credit ratings. Although we cannot provide any guarantees that these technologies provide an easy route to financial riches, we hope this book will spark new ideas in the minds of readers to encourage them to undertake their own work in the fascinating nexus of computer science and finance.

This book is aimed at two audiences: those in the finance community who wish to learn about advances in biologically inspired computing and how these advances can be applied to financial modelling; and those in the computer science community who wish to gain insight into the domain of financial modelling and trading system design. Strong emphasis is placed in this book on evolutionary methodologies, particularly *Grammatical Evolution* [174]. This book is also suitable for use on advanced undergraduate or post-graduate courses, on quantitative finance or computational intelligence. No prior knowledge of either BIAs or financial prediction is assumed.

## 1.1 Biologically Inspired Algorithms

Biological systems are a notable source of inspiration for the design of optimisation and classification algorithms, and all of the methodologies in this book have their metaphorical roots in models of biological and social processes. These processes are as diverse as the operation of the central nervous system, biological evolution, the mapping of genes to proteins, the human immune system, and models of social interaction amongst organisms. BIAs do not seek to perfectly imitate the complex workings of these systems, rather they draw metaphorical inspiration from them to create mathematical algorithms which can be used in an attempt to solve hard, real-world problems, such as modelling financial markets. Figure 1.1 provides a broad taxonomy of some of the primary methodologies discussed in this book. A vast number of hybrid models which combine elements from more than one of these methodologies can also be constructed.

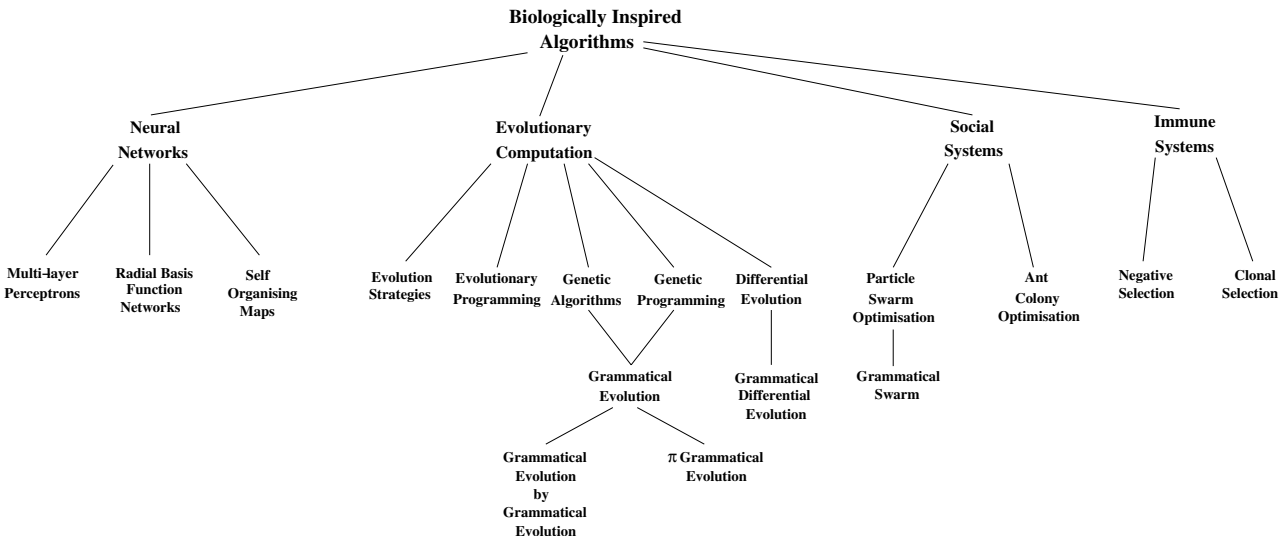
It is not possible to undertake a complete discussion of all of these in a single text, and we concentrate on neural network and evolutionary algorithms, while providing an introduction to BIA technologies drawn from social and immune metaphors. A brief overview of some of these technologies is provided in the following paragraphs, with a more detailed discussion of them being provided in later chapters.

### 1.1.1 Artificial Neural Networks

Artificial neural networks (NNs) is a modelling methodology whose inspiration arises loosely from a simplified model of the workings of the human brain. Both learn from their environment and encode this learning by altering the connections between individual processing elements, neurons in the case of the human brain, nodes in the case of NNs. NNs can be used to construct models for the purposes of prediction, classification and clustering. NNs are a non-parametric modelling tool, as the model is developed directly from the data.

### 1.1.2 Evolutionary Computation

Evolutionary algorithms draw inspiration from the processes of biological evolution to *breed* solutions to problems. These problems may be as diverse as determining the coefficients for a non-linear regression model, or determining the components of a financial trading system. The algorithm commences by creating an initial population of potential solutions, and these are iteratively improved over many ‘generations’. In successive iterations of the algorithm, fitness-based selection takes place within the population of solutions. Better solutions are preferentially selected for survival into the next generation of solutions, with diversity being introduced in the selected solutions in an attempt to uncover even better solutions over multiple generations. BIAs that



**Fig. 1.1.** A taxonomy of the biologically inspired algorithms that are discussed in this book

employ an evolutionary approach include genetic algorithms (GAs), genetic programming (GP), evolutionary strategies (ES) and evolutionary programming (EP).

A significant recent addition to BIA methodologies is grammatical evolution (GE), an evolutionary automatic programming methodology, which, for example, can be used to evolve rule sets or financial trading systems. GE incorporates a *grammar* which governs the creation of these rule sets. The idea of a grammar is inspired by the biological process of the mapping of genes to proteins.

### 1.1.3 Social Systems

The social models considered in this book are drawn from a *swarm* metaphor. Two popular variants of swarm models exist, those inspired by the flocking behaviour of birds and fish, and those inspired by the behaviour of social insects such as ants. The essence of these systems is that they exhibit flexibility, robustness, self-organisation, and communication between individual members of the population. The swarm metaphor has been used to design algorithms which can solve difficult problems by creating a population of problem-solvers, and allowing these to communicate their relative success in solving the problem to each other. Higher-performing individuals attract the attention of others, who test variants on their problem-solving strategy in an attempt to improve it.

### 1.1.4 Artificial Immune Systems

The human immune system is a highly complex system, comprised of an intricate network of specialised tissues, organs, cells and chemical molecules. The capabilities of the natural immune system are to recognise, destroy and remember an almost unlimited number of foreign bodies, and also to protect the organism from misbehaving cells in the body. To assist in protecting the organism, the immune system has the capability to distinguish between *self*, and *non-self*. Artificial immune systems (AIS) draw inspiration from the workings of the natural immune system to develop algorithms for optimisation and classification. Practical applications of AIS models to pattern-recognition tasks include the identification of potentially fraudulent credit card transactions, the identification of the ‘state’ of the stock market, and the identification of financially at-risk companies.

## 1.2 Computer Trading on Financial Markets

Computerised or automated trading on financial markets is not a new phenomenon. Computers have been used for *program trading* for many years. In



program trading, the object is usually to uncover and eliminate anomalies between financial derivatives and the underlying financial assets which make up those derivatives.<sup>1</sup> A typical example of program trading is *index arbitrage* which involves the automated purchase or sale of a basket of stocks which make up a market index, in conjunction with the simultaneous sale or purchase of a derivative product such as stock index futures, in order to profit from the price difference between the basket and the derivative product. In theory the transaction generates risk-free returns, but in practice it relies on estimates of dividend income from companies, an estimate of the rate of return available on invested dividends, and the ability of the computer to make the purchases/sales at the prices which produced the arbitrage opportunity. Program trading accounts for a considerable portion of trading on major stock exchanges. For example, it is estimated that program trading volume accounted for approximately 50.6% of the total trading volume on the New York Stock Exchange (NYSE) in 2004 [163].

A second, less publicised use of computers is to construct trading systems which assume trading risk in the search for superior, risk-adjusted, returns. These systems are the focus of interest of several of the case studies in this book.

### 1.3 Challenges in the Modelling of Financial Markets

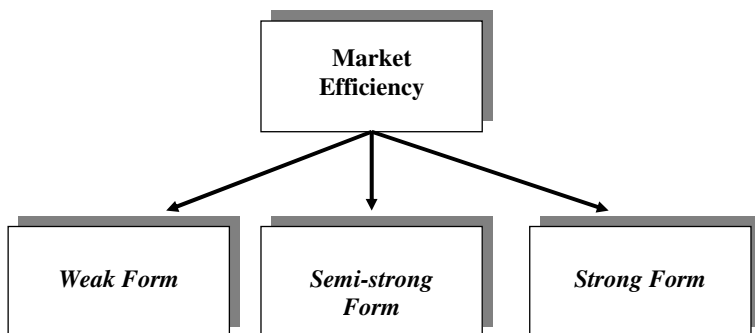
Modelling of financial markets is challenging for several reasons. Many factors plausibly impact on financial markets including interest rates, exchange rates, and the rate of economic growth. We have no hard theory as to how exactly these factors effect prices of financial assets, partly because the effects are complex, non-linear and time-lagged. For example, a change in interest rates may impact on the foreign exchange rate for a currency, in turn effecting the level of imports and exports into and from that country. Another difficulty that arises in financial modelling is that unlike the modelling of physical systems we cannot conduct controlled experiments. Only one sample path through time is available for our examination, as we only have one history of market events. Additionally, some factors which can effect financial markets are inherently unpredictable such as earthquakes, the weather, or political events. Taken together, these difficulties imply that our ability to predict market movements will always be imperfect.

---

<sup>1</sup>A derivative is a financial instrument whose value is based on that of another financial instrument such as a share. For example, investor A may sell an option on a share to investor B. This option gives investor B the right to buy (or sell) that share to investor A, at a specified price for a specified time. As the value of the underlying share changes, the value of the financial derivative (the option) will also change.

### 1.3.1 Do Prices Follow a Random Walk?

The very attempt at modelling financial markets for profit meets with the scorn of some financial economists. Two pillars of traditional financial economics are that market prices of financial assets follow a *random walk* and that markets are *efficient*. One of the earliest studies suggesting that prices in markets might follow a random walk was undertaken by [123]. The traditional definition of a random walk is a process in which the changes from one time period to the next are independent of each other, and are identically distributed. If prices in financial markets did follow a random walk, this would imply that the size and direction of a past change in price provides no insight into the size and direction of the next change in price of that financial asset. In other words, there would be no auto-correlation in the time-series of prices from a financial market.



**Fig. 1.2.** Three forms of market efficiency

Closely related to the concept of a random walk (and often confused with it) is the efficient market hypothesis (EMH). Although a random walk in share prices could arise for a variety of reasons, it is consistent with a proposition that current prices fully reflect the market's aggregate assessment of any existing information which could impact on the price of a financial asset. If a market is informationally efficient, in that all information is impounded accurately and instantly into prices once it becomes available, then there is no scope to make excess returns from trading on such information. The more information-efficient a market is, the more random the sequence of price changes that it will produce, as prices will only alter when new information emerges. As the nature of new information is by definition unpredictable, a share's price is equally likely to rise or fall in the future. Three versions of market efficiency were posited [73] (Fig. 1.2).

Under weak form efficiency it is considered that the price of a share at any point in time reflects *all* the information contained in its price history. This would imply that excess risk-adjusted returns cannot be obtained by attempting to construct a model which uses information on past share prices, or past transaction volumes, to predict future share prices. The semi-strong form of efficiency suggests that a share price at any point in time reflects all publicly available information which could impact on the share's price. The strong form of market efficiency implies that share prices include all information relevant to the price of a share, including both public and private (non-public) information.

### 1.3.2 Attack of the Anomalies

If the semi-strong form of the EMH was correct, there would be no scope to construct a model of a financial market using publicly available information, which would generate excess risk-adjusted returns. In spite of the initial research which lent broad support to the EMH, there is a growing body of research in more recent times which suggests that subtle patterns do exist in time-series of financial asset prices, and that these prices do not follow a random walk [143, 144]. Among the anomalies that have been noted are the existence of serial correlation in weekly and monthly stock returns, particularly for small capitalisation (small company) stocks. Generally, three patterns of serial correlation in stock returns are recognised: short-term reversals (looking at returns over a few weeks), medium-term inertia, and longer-term reversals.

A considerable body of empirical evidence suggests that short-run volatility in share returns is clustered. A large change in price tends to be followed by another large change in price, but the direction of this change is difficult to predict. In other words, prices tend to be volatile when they have just been volatile, leading to patterns of short-run price reversals. Under medium-term inertia, good (or bad) performance of a stock over 3-12 months is typically indicative of continued good (or bad) performance in the next few months. This could provide scope for the implementation of *momentum* investment strategies, where investors seek to buy (sell) shares which have recently trended upwards (downwards). Over longer time periods (3-5 years), there is evidence of negative serial correlation in share returns, suggesting that stocks that have performed well over the last several years are more likely to under-perform in the future. This lends support to the common idea of *contrarian* or *value* investment strategies where investors buy out-of-favour stocks (those whose share price has underperformed that of their peers in recent years). A posited explanation for this negative serial correlation is the *over-reaction hypothesis* that investors are subject to waves of optimism and pessimism which cause prices to swing temporarily away from their underlying value for individual firms or whole sectors [53, 63]. Other examples of asset-price anomalies are reported in [37] and [54].

Interpretation of the results of studies reporting anomalies has been controversial [74, 75], but they are consistent with a hypothesis that market efficiency is a relative term. Under this premise, as market participants uncover new information processing mechanisms (such as BIAs), market efficiency is enhanced as market participants apply the new information processing methodology. It is plausible that novel, powerful computational techniques which can uncover new price-relevant information could prove profitable. However markets represent a competitive, adaptive environment and are likely to rapidly impound available information into asset prices. Just as \$100 bills do not last long on the sidewalk, traders who constitute markets have a vested interest in searching for and exploiting any edge which could lead to profit. Financial modelling in such an environment can be compared to an arms race whereby each player rapidly cancels out any advance of another. The advantage offered by a new modelling technique is therefore likely to be short-lived.

There are close parallels between the challenges of the environment of financial markets, and those from which several of the biologically inspired algorithms discussed in this book are drawn. Evolutionary algorithms, swarm algorithms and immune algorithms are drawn from environments where, just as in financial markets, there is continual adaptation and where there is competition for resources between individuals.

## 1.4 Linear Models

The goal in modelling a system is typically to gain insight into its behaviour, to determine which factors impact on the output of the system, and to determine how influential each of these factors is. A second goal is to enable prediction of the future output of the system under different conditions. A simple linear model has the general form:

$$Y = \alpha + \beta_1 X_1 + \dots + \beta_n X_n \quad (1.1)$$

where  $Y$  is the dependent variable,  $X_1, \dots, X_n$  are independent (or explanatory) variables (in a simple model there may be only one independent variable, in a multiple regression model there will be several),  $\beta_1, \dots, \beta_n$  are regression coefficients, and  $\alpha$  is a constant which allows the model to produce a value for  $Y$  even if all the dependent variables have a zero value.

In constructing a linear model, the first step is typically to posit a *cause-and-effect* relationship based on prior theory or intuition between the dependent variable and one or more explanatory variables (or inputs). In other words: ‘I think  $x$  and  $y$  impact on the value of  $z$ ’. Two questions then spring to mind:

- i. Is the assumption that  $x$  and  $y$  effect  $z$  supported by empirical evidence?
- ii. How great is the effect of  $x$  and  $y$  respectively on  $z$ ?