Agent inactivity in the Evolutionary Minority Game

by

Michael D. Mitton

B.Sc. (Physics) The University of British Columbia, 1996

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Master of Science

in

THE FACULTY OF GRADUATE STUDIES

(Department of Physics and Astronomy)

We accept this thesis as conforming
to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA

April 2001

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the head of my department or by his or her representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Department of Physics and Astronomy

The University of British Columbia
Vancouver, Canada

Date April 25, 2001

DE-6 (2/88)

Abstract

In this work I have developed a model, based on the Evolutionary Minority Game (EMG), to examine the effect of agent inactivity within a group of agents competing for a limited resource, but governed by both supply-and-demand and a minority rule. The structure of the inactivity mechanism has been modeled after the strategy preference parameter, p, of the EMG. A parameter has also been introduced that models, in a very simple way, the effect of inflation on the system in order to motivate the agents to play the game.

The behaviour of the model has been examined with the use of numerical simulations over its entire phase space. The results focus on two aspects of the model: 1. how an agent's performance depends on his activity level; and 2. how the properties of the EMG are affected when agents are given the option of not playing the game. Results of these simulations demonstrate that an agent's performance is strongly dependent on his level of activity. Specifically, it has been shown that an agent's optimal level of activity is dependent on his strategy preference and, moreover, that this optimal activity level undergoes a first-order transition from a phase of optimal inactivity to a phase of optimal activity as the inflationary force is increased. Even though an agent's activity decision has been modeled as a process independent of his decision to "buy"

or "sell" during a given round, results of simulations indicate that correlations between the two decisions have emerged from the collective dynamics of the group.

A theory of the model has also been developed. The formulation of the theory is based on a mean-field approach in which the actions of a single agent are considered within a background field produced by the remainder of the group. It has been found that the results of the theory agree with those of the numerical simulations very well and appear to become exact in the thermodynamic limit. Furthermore, the discrepancy between theory and simulation at finite N indicate a possible breakdown of the non-local nature of the inter-agent interactions built-in to the model.

Contents

List of Figures

Acknowledgements

Firstly, I would like to thank Dr. Birger Bergersen for taking me on as his student, providing me with access to the new and exciting field of "econophysics", and allowing me the freedom to choose and develop this thesis topic.

I would also like to thank Dr. Alan Kraus for providing me with helpful suggestions and keeping my feet on the ground with regards to the financial validity of the Minority Game model.

Finally, I would like to thank Dr. Hendrik Blok for his constant input throughout the duration of this work and his patience with the many, semi-prepared, late-night "reference" visits that kept him from a warm dinner. ps - a post-doc *does* hold weight

MICHAEL D. MITTON

The University of British Columbia

April 2001

Chapter 1

Introduction

1.1 Financial asset prices under speculation

A financial market is a structured organization designed to centralize, and therefore facilitate the trade of similar types of financial *assets*, such as stocks, bonds, commodities, and derivative securities. The specific mechanism by which a trade occurs varies from market to market, but in general there are traders (which I shall term *agents*) who submit *bids* to buy an asset at a certain price and therefore create *demand* for that asset. Opposing this flow of demand is the *supply* created by agents who submit *offers* to sell the asset. In very general terms, a trade is then executed in one of two ways [3]: in a *centralized* market, a central third-party, or *market maker*, weighs the total supply and demand and sets a price to clear the market; in a *decentralized* market, individual agents simply meet face-to-face and agree on a price at which to exchange the asset.

From the above description it might seem that the process of asset price formation is relatively simple: agents simply determine the value of the asset by the value of the underlying (ie. the value of one non-dividend-paying stock of a company would simply be the value of the company divided by the number of stocks outstanding) and buy (sell) when its price is lower (higher) than this value. In this scenario, only the change of exogenous factors, such as interest rates, the Gross Domestic Product, and foreign exchange rates, or the release of news concerning the underlying should cause a change in the asset's price (although it should be noted that even a change in the value of the underlying will not necessarily lead to trading of the asset). When none of these factors are present each agent should hold the same opinion for the value of the asset and no trading should occur.

In reality, we know that trades of most assets occur at a high frequency not only during the release of news, but also in between these events. This observation is indicative of a very complicated mechanism of asset price formation and is mainly a result of two factors. First, not all agents have the same opinion about the "true" value of an asset and about how it will change in the future. These individual beliefs can result from an agent's risk preference, social conscience, or simple irrationalities; since they are not necessarily tied to anything tangible, these *speculative* effects are typically very hard to quantify. Second, the vast communication structure connecting agents within and across markets further complicates the dynamical behaviour of speculative asset pricing. With the progress of information technology in recent years, the physical barriers within financial markets are crumbling and the network of interactions between agents is increasing rapidly resulting in trading decisions that are now rarely made independent of other agents' opinions and decisions.

Empirical properties of speculative asset prices

The above mechanisms produce very interesting properties observed in asset price time series. Let X_t be the asset's price at time t and $r_{t,\Delta t} = (X_t - X_{t-\Delta t})/X_{t-\Delta t}$ be its *rate of return* at time t and over the time interval Δt. It has been well documented [4–6, 9, 14, 15, 17, 18] that the probability density function, $f_{r_{t,\Delta t}}(x) = Prob[r_{t,\Delta t} = x]$, is distinctly *leptokurtic* for time-scales, $\Delta t \leq 1$ month; as $\Delta t \to \infty$, this leptokurtic form converges to a Gaussian.

It has also been shown that interesting correlations exist between successive asset price movements [4–6, 11–14, 16]. Let $\hat{E}[\,]$ be the expectation operator and $\mu_x \equiv \hat{E}[x_t]$ be the mean of some time series x_t. Specifically, the autocorrelation function of returns

$$\rho_r(t_1, t_2) \equiv \frac{\hat{E}[(r_{t_1,\Delta t} - \mu_r)(r_{t_2,\Delta t} - \mu_r)]}{\sqrt{\hat{E}[(r_{t_1,\Delta t} - \mu_r)^2]}\sqrt{\hat{E}[(r_{t_2,\Delta t} - \mu_r)^2]}} \tag{1.1}$$

for an asset within a well-established and highly liquid market decays rapidly (specifically, with half-life $\tau \approx 4$ min) to the level of system noise within approximately 30 minutes. This result is somewhat expected since it reflects the time for new information to be reflected in the asset price. More interestingly, though, is the very slow decay of the autocorrelation function of volatility (where the volatility, or standard deviation, $v_{\delta t}(t)$, is defined as the RMS return over a time, $\delta t \gg \Delta t$)

$$\rho_v(t_1, t_2) \equiv \frac{\hat{E}[(v_{\delta t}(t_1) - \mu_v)(v_{\delta t}(t_2) - \mu_v)]}{\sqrt{\hat{E}[(v_{\delta t}(t_1) - \mu_v)^2]}\sqrt{\hat{E}[(v_{\delta t}(t_2) - \mu_v)^2]}} \tag{1.2}$$

Even for $t_2 - t_1 >$ several years, $\rho_v(t_1, t_2)$ is still significantly above the level of noise. This effect, known as *volatility clustering*, simply means that an asset price movement tends to be followed by one of a similar size; it has received much attention since it shows that there are long-range correlations within $|v_{\delta t}(t)|$ and it might therefore be possible to *predict* the asset price's future movements.

1.2 The physics of speculative asset pricing

1.2.1 Why should physicists study finance?

Although numerous variations have been developed, to date, no model has been put forward that is able to accurately generate all of the known empirical characteristics of speculative asset price time series. Historically, the majority of this effort has come from the financial economics community; several popular examples of these models are Mandlebrot's Stable Paretian Hypothesis [1], Clark's Subordinated Stochastic Process (SSP) model [27], and the (Generalized) Autoregressive Conditional Heteroscedasticity ((G)ARCH) class of models (see, for example [28–30]). While it *is* a problem that none of these *econometric* models are able to fully capture all of the stylized facts of the empirical data, much progress has been made since their introduction and it is feasible that this goal will be met in the future.

Of a much greater concern is the fact that econometric models typically give very little cause-and-effect information about *how* the underlying market mechanisms generate the empirically observed stylized facts. Without knowledge of how the model produces the desired asset price dynamics our use of it would be "blind" and it would not be known whether its assumptions were valid from one time to the next. As physicists, it is our natural inclination to seek out the fundamental cause to a problem, which, in this case, should also provide us with both practical and academic benefits.

Over the past few years, physicists have shown a lot of interest in speculative asset pricing and financial systems in general. The idea behind the physics of finance, or *econophysics*, is that financial markets act in a very similar way to certain statistical physical systems and therefore should be amenable to the tools developed to treat them. There is no question that researchers in the fields of finance and economics have been studying these problems for some time and have had a great deal of success. The goal of the econophysics field is not to replace this work, but rather to provide an approach to complement it.

1.2.2 Complexity in finance

Along with the goal of producing a model that gives insight into the cause of an asset price's behaviour, the motivation to study finance has, in large part, come from the realization that financial systems (ie. markets) are in many ways similar to the general class of "complex" physical systems such as frustrated magnetic systems, neurological networks, and surface growth.

Structural similarities

The most obvious similarity between a financial market and one of these physical systems is in their underlying structures. Complex systems always possess a large number of interacting subunits. While the form of this subunit is human in the case of a financial system (as opposed to an atom or molecule typical of physical systems), markets are *very* large structures with an even larger communication network connecting one another making the form of the macroscopic systems very similar.

There are, however, two significant structural differences between the two classes of systems. First, in a financial market each subunit is a living organism and therefore a complex system in its own right. While the internal structure can be very complicated for some physical systems (ie. the individual phospholipids of a biological membrane), it pales in comparison to that of a human being. Second, in a physical system, the laws governing the interaction and evolution of the microscopic components are well-defined fundamental forces; in a financial system, however, many of these forces are the result of poorly understood cognitive factors. It is interesting to wonder how well techniques designed for systems with particle-like subunits and which evolve under physical laws will adapt to social systems governed by less well-defined forces.

Empirical similarities

As the number of subunits within a macroscopic system is increased, it is intuitive to reason that the complexity of its behaviour will increase proportionately. It is a well documented fact that this is not the case. As the system's size increases, regular and distinctive macroscopic properties actually emerge from the complex microscopic dynamics.

To illustrate these properties, consider for example the magnetization, M, of a macroscopic magnetic material [20] at a temperature, T immersed in an external magnetic field, H. If $H = 0$ and the temperature is increased, it is found that at some critical temperature, T_c, the system enters a *critical state* in which: 1. correlations between individual "spins" of the magnet become very long-range; 2. power law behaviour (indicating scale-invariance) is observed; and 3. the magnetization undergoes a transition and falls sharply to zero.

Rescaling the temperature so that $\tau = (T - T_c)/T_c$ and then measuring M vs. τ for several values of H yields a very interesting result. It is found that, if M is rescaled by some power of H, and τ by some different power of H, plots of the rescaled M versus the rescaled τ all collapse onto a single curve. This *scaling law* (the specific values of the scaling exponents for a given system) and *data collapse* for this magnetic system are traits common to complex physical systems.

Another interesting characteristic of such systems is termed *universality*. To demonstrate this effect, imagine the above $M = M(H, \tau)$ measurements are repeated on several different types of magnetic materials. From the measurements it is found that, ignoring the necessary material constants, the same scaling law holds for each material; it is said that all of these systems belong to the same *universality class* which means that, as the critical state is approached, their macroscopic properties behave in a similar manner.

Since financial markets posses a similar structure to traditional complex systems it is interesting to wonder whether they also share the same empirical characteristics such as scaling behaviour and universality. Many studies have been performed along these lines [1,2,7–9,16–24] and two key results have been found. First, the *cumulative* distribution of returns, $F_{r_{t,\Delta t}}(x) = Prob[r_{t,\Delta t} \geq x]$ is found to asymptotically decay as a power law with exponent, $\alpha \approx 3$. When the rate of return is scaled with the appropriate exponent, $F_{r_{t,\Delta t}}(x)$ is also found [1, 2, 4, 5, 7–9, 16–19] to be stable for time scales $\Delta t \leq 20$ days (for individual stocks). Second, it is also now well known that the asymptotic power law decay exponent, $\alpha \approx 3$ and the appropriate scaling exponent hold for individual stocks, whole-market indices, and even across various equity markets suggesting universality of the price formation process in these markets.

Finally, the fact that long-range autocorrelations exist in an asset price's volatility time series and its cumulative distribution of returns exhibits power law behaviour suggest that a financial market exists normally near its critical state. This behaviour differs from that observed in the magnetic example above in that the magnet will only reach a critical state if its parameters, τ and H, are *tuned* to their critical values. One possible explanation for this phenomenon is the *Self-Organized Criticality* hypothesis proposed by Bak et al. [50] in which a state of dynamical equilibrium is reached once the system has evolved to the point where it can no longer propagate a perturbation over the correlation length of the system.

1.2.3 Physical models of asset price dynamics

The evidence is quite strong that the mechanism underlying the process of speculative asset price formation is very similar to that of various "complex" physical systems. The question is now *How do we use the standard tools of statistical mechanics to develop a useful model of speculative asset prices?*

The number of different physical models of speculative asset price dynamics is very large (see for example [32–49]) and to give a description of a "general" method would do none of them any justice. Instead I shall briefly describe the structure (and its motivation) of *one* of these models which I believe to be a good representation of the efforts of this approach.

The model I shall discuss is the Cont-Bouchaud (CB) model [31]. In the CB model, correlations are introduced by incorporating a communication structure between the N individual agents who are situated at the vertices of an N-dimensional random graph. Each agent is allowed to randomly establish binary links with other agents and, through successive links between existing clusters (two or more agents linked together), large clusters can form; in this way the communication structure is analogous to bond percolation in an infinite-dimensional space.

Within a single cluster it is assumed that all agents hold the same position towards the asset (ie. buy, sell, or hold) and so the price change (which is assumed to be proportional to the the excess demand) is simply proportional to the cluster weighted total demand. A key aspect of the model is an agent's ability to form coalitions; in the CB model this is controlled by a constant, user-set parameter. For a range of values of this parameter, the model generates $f_{r_{t,\Delta t}}(x)$ in a manner that is very similar to what is observed in real markets. Specifically, it predicts both exponentially truncated tails (with $\alpha > 2$) and $\alpha = 1.5$ in the region of small-to-moderate $|r_{\Delta t}|$ which is very close to what is observed empirically. Finally, it is found that, as the clustering parameter approaches some critical value, very large coalitions form and a market crash occurs.

As well as generating quantitatively accurate data, the CB model is useful for showing *how* a specific mechanism (in this case the inter-agent communication structure) can result in the leptokurtic form of $f_{r_{t,\Delta t}}(x)$.

1.3 The Minority Game

1.3.1 The model

The Minority Game (MG) was developed by Challet and Zhang [55] as a simplification of Arthur's El-Farol Bar-attendance problem [54]. The model is designed to examine the behaviour of a group of individuals who all compete for a limited resource and are governed by the rule of supply-and-demand.

The MG consists of a non-spatially oriented, odd number of agents who repeatedly attempt to out-guess one another and earn a reward. On each round of the game, the agents take one of two possible positions; with foresight in hand I will label these *buy* and *sell*. After the agents have made their choices, the size of each position is tallied; those agents who chose the position taken by the minority of agents are rewarded with $+1$ point, while those who took the opposite (majority) position are penalized -1 point.

To make his decision, an agent has a set of *strategies.* Each strategy predicts a specific position for the agent to take given each unique *history* of the game. The history consists of a certain number of the most recent winning positions of the game and is common to all agents within the group. Finally, in order to incorporate adaptability into the game, each agent rates the performance of his strategy(ies) and makes his strategy selection based on this performance.

Many details were omitted from the above description of the MG since there are several formulations of the model which are all slightly different. In the original MG (which I will refer to as the *Standard* Minority Game), each agent possesses a finite and (possibly) unique set of strategies randomly drawn from some large simplex space. A virtual score is kept for each strategy and a virtual point is awarded on any round in which the strategy would have predicted the minority position. In the Standard MG agents use the simple strategy selection rule of choosing their strategy on each round with the highest virtual point score.

The Standard MG is useful for examining global properties (such as system efficiency) of the group of agents within a MG framework. In a slightly generalized version of the Standard MG, agents have the ability to select *any* of their strategies (not just the one with the highest virtual score) according to the Boltzmann weight of their virtual scores. Because of the ever-changing dynamic of the game and since the agents' strategy sets can overlap, playing what is thought to be a "sub-optimal" strategy may in fact be beneficial if enough of the other agents play their similar "optimal" strategies. In this way, this *Thermal Minority Game* is useful for examining whether-or-not using a sub-optimal strategy is actually irrational in this complex environment.

1.3.2 The Evolutionary Minority Game

While the agents within the Standard and Thermal MG models are adaptive (the agents can select from a variety of strategies depending on the current conditions), they are not strictly evolutionary since both the population of agents and each agent's strategy set is fixed throughout the duration of the game. In reality, we would expect that an underperforming agent or strategy would be replaced with one that appears to be better equipped to compete within the group dynamics.

Another popular version of the Standard MG, designed to include evolutionary effects, is the *Evolutionary Minority Game* [70]. In this model there are only two strategies available to the agents; one which contains the minority positions the last time each of the possible histories was encountered and the other with the opposite positions (ie. the majority positions); in this way the two strategies are dynamic and change constantly throughout the game. Each agent then has a unique parameter (called his *gene value*) that predicts which of the two strategies he will play. Evolution is introduced by removing from the population agents whose performance falls below a certain level and replacing them with agents having different gene values.

1.3.3 The MG as a realistic market model

As should be clear from the above description of the MG, the structure of the game is *very* simplistic. One might ask (justifiably) why the MG should be included in the discussion of speculative asset price models. To be perfectly honest, none of the versions of the MG presented thus far are remotely close to being accurate representations of a real financial market. Some problems with using the MG as a realistic model are:

- Each agent is constrained to be active on every round of the game.
- The size of the group is fixed.
- There is no external information present in the model.
- The binary payoff function ($+1$ for taking the minority position and -1 for taking the majority position)is overly simplistic.

While these details definitely restrict the applicability of the Minority Game to financial systems, they are simply structural components of the model and can be removed and replaced (fairly easily) by more realistic constructions. For instance, in a real market the total sales and purchases of an asset must equal; this fact implies that the total wealth is conserved within the group trading in the asset (neglecting commissions, spreads, etc...). The binary payoff function of the MG is inaccurate since it does not preserve the system's wealth. Given there are n_{min} agents in the minority group and n_{maj} majority agents, then one simple method of incorporating this conservation law while maintaining the model's basic minority structure would be to give a reward of $+1$ point to each of the minority agents and a penalty of $-n_{min}/n_{maj}$ points to each of the majority agents.

Of a far greater concern are the components of the MG that are fundamentally incompatible with a financial market; these include:

- In the MG there is no motivating factor for the agents to play the game. Even though speculation is a crucial component of a real market, for the agents to trade they must at least believe that the asset has *some* intrinsic value. In the MG this does not exist; the agents simply play the game because we force them to.

- The MG has no market making mechanism; agents simply "buy" and "sell" the asset. It is assumed that there is an unlimited quantity of the asset and all of those sold are then bought and the quantity of those to be bought are present in the market to begin with.

- The MG neglects the effect of an agent's actions on his long-term performance. If there are more buyers than sellers during a round then we should expect the price to rise. While selling the asset during this round will earn an agent a good short-term profit, that agent will then own less of the asset when the price rises which is not beneficial in the long run. The MG essentially models how agents compete under supply-and-demand pressures, but neglects any effects due to the arrival of external information (which result in the long-term movements of the asset price).

With its construction, the MG was not intended to act as a model of a financial market; it was simply meant to model how a group of heterogeneous individuals behave when competing for a limited resource and governed by a minority rule. For instance, in its basic form the MG acts as a fairly realistic model of: 1. commuters competing for the less crowded of two possible routes; or 2. foraging animals searching for a limited food source. Because of the reasons listed above the MG is too simplistic to be considered as a realistic market model.

However, some researchers [56, 57, 62, 66, 69] have discussed the MG in terms of a market model. Neglecting the above problems, the MG *does* have features that make it attractive for this use.

Minority rule

It was stated above that the movements of an asset's price usually follow the actions of the majority. While this effect is central in determining the long-term performance of an agent, under certain circumstances the *minority rule* effect is present and significant. Imagine the situation of a group of fund managers competing for profits through ownership of two stocks, x_1 and x_2. Assume it is a commonly held belief that both stocks will increase in value, but that x_1 will increase at a slightly higher rate than x_2. If we further assume that : 1. there are equal amounts available of both x_1 and x_2; 2. each manger can only invest in *one* of x_1 or x_2; 3. the total quantity of each asset is shared equally among those who buy it: and 4. all resources not used to purchase x_1 or x_2 are held in the form of some more secure, but lower return asset such as a government bond, then it might seem logical for each manger to ignore x_2 and buy as much of x_1 as possible. But if this happens then none of the managers will do very well since they will have to share the profits amongst a very large group. If the commonly held belief holds true and x_2 earns slightly less per share, then a small (minority) group of the managers will outperform those in the majority group if they buy x_2 since they will be able to buy more and earn a greater profit.

Inductive reasoning

Traditional financial economics has always assumed the use of a *deductive* agent and the Efficient Market Hypothesis (EMH); given any economic situation, an agent makes a logical decision based upon a perfectly defined environmental state in which all information about an asset is reflected in its current price. For instance, within an equity market, such a perfectly rational agent would weigh all information (such as the stock's price history and the "true" value of the underlying company) and trade according to the relationship between the current price and this information.

In real financial systems these conditions are often very complex and sometimes ill-defined; assuming that an agent can decipher all of the information perfectly is not reasonable. It is known [54], however, that humans are very good at pattern matching and therefore reasonable to expect that, under complex circumstances, we will use *inductive* reasoning and base our decisions on our past experience.

As opposed to standard models of speculative asset pricing, the dynamics of the Minority Game are formed through the use of inductive reasoning. On any round, an agent selects a position based on the current history of the game and his personal experience with that history. Using this experience, the agent can then adapt (in a *Lamarkian* way in the Standard and Thermal MG) or evolve (in the *genetic* sense of the EMG) in order to be more competitive within the ever-changing dynamic of the game.

The Minority Game approach

I believe the Minority Game is a convenient starting point to develop a market model which is capable of generating empirically accurate results. Although many details of the structure of a financial market are not present in the model, it *does* provide an attractive model (with its use of inductive reasoning) of how a group of agents compete for a limited resource when governed by a minority rule with dynamics resembling the process of price formation in many situations.

The idea central to the MG approach for market modeling is to first understand the dynamics underlying the price formation process and then to build a more realistic model upon this base. There are two steps behind this approach.

1. Characterize the basic MG model.
2. Using this model as a base, successively remove its simplifying assumptions and gradually build-up the more realistic model.

Actually, the simplicity of the MG permits this approach; using such a simple model provides a tractable base that is relatively easy to treat analytically.

While this approach is time-consuming, it has several benefits. First, beginning with a simple base and then building the model up component-by-component should allow us to determine which components (market mechanisms) are necessary for the model to produce results in agreement with those observed empirically. Second, it is also hoped that this type of approach will allow us to understand the specific effect(s) of each component and provide us with a cause-and-effect map between these microscopic components and the macroscopic dynamics of the asset price. Finally, from the vast number of current speculative asset pricing models it is clear that the structure of these components can be modeled in a wide variety of ways (for example the payoff function can be modeled as a simple step function as in the basic MG, or it could be modeled (slightly) more realistically as a fixed level of resource split evenly between each member of the minority group); with this bottom-up approach we should be able to gain insight into how the specific structure of each component affects the asset price dynamics.

Progress to date

Since its introduction, much work has been done on the Minority Game (see, for example [55–72]). This has resulted in both a thorough understanding of the dynamical behaviour of the models and in accurate analytical theories of both the Standard MG [61] and the EMG [71].

Many extensions of the models have also been studied. First, a more realistic payoff function (similar to the one suggested in 1.3.3) that conserves the wealth of the system has been studied for the Standard MG [66]. It has been found that the properties of the model are largely independent of the specific structure of the payoff function. In addition, the effect of various strategy distributions on both the Standard [63] and Evolutionary [72] MG models has been investigated as well as the size and structure of the agent's strategy space [67, 68].

Along with many other studies investigating single extensions of the basic MG model, there have also already been two attempts at large-scale extensions of the model with the goal of creating price dynamics with empirically accurate features. Removing only 3 or 4 of the simplifying assumptions of the basic MG, these papers [57, 62] have already produced results that are in good qualitative agreement with both the empirical distribution of returns, $f_{r_{t,\Delta t}}(x)$, and the autocorrelation of returns, $\rho_r(t_1, t_2)$.

1.4 Focus of thesis

1.4.1 Motivation for research

Apart from a brief treatment in [62], the effect of agent inactivity on the Minority Game model has largely been ignored. The lack of attention to this phenomenon is puzzling; from both day-to-day experience and the results of academic research, inactivity within a group of competing individuals seems to be an important component of the price formation mechanism. In their paper [31], Cont and Bouchaud state that allowing the agents to sit out at various times is crucial to obtaining the leptokurtic $f_{r_{t,\Delta t}}(x)$ and volatility clustering characteristic of empirical data. From real-world experience it should be obvious that: 1. the number of people trading in any particular asset is constantly changing; and 2. investors are not forced to trade continuously and so their trading frequency will also vary in time.

1.4.2 Present study

In this thesis I will examine one very simple model, called the Variable Activity Evolutionary Minority Game (VAEMG), of agent inactivity within the Minority Game framework. In keeping with the ultimate goal of the MG I will attempt to characterize the model as completely as possible and, in doing, will focus on two aspects: 1. how the performance of an arbitrary agent evolves when he is given a variable activity level; and 2. how the dynamics of the game are affected under these same circumstances.

In analogy with the gene value of the EMG, each agent in the VAEMG will have his own activity parameter which is simply the probability that the agent is active on any given round. The structure of this inactivity mechanism is intentionally simplistic so as to maintain tractability in the model. At the moment I am simply concerned with how the *presence* of an inactive state affects the behaviour of the game; it will be left to a later study to examine more complex and realistic inactivity mechanisms.

The play of the VAEMG is split into two steps. In the first, each agent "decides" whether to be active or not. If an agent chooses to participate in the round then he simply plays an EMG-type game with all of the other active agents in the second step. To motivate the agents to be active, I have included a small penalty to an agent's score for each round he is inactive. When an agent's score falls below some predetermined level that agent is killed-off and replaced by one with new activity and gene values that are hopefully better able to compete within the group dynamic.

In this thesis I will only examine the effect of this inactivity mechanism within the framework of the EMG model. This choice has been made since the EMG is ideal for examining the properties of an individual agent and I am interested in how this performance will be affected when the agent has the ability to choose whether or not to participate in the game.

1.5 Organization of thesis

The remainder of this study will be as follows: in chapter 2, I shall develop the formal definition of the EMG model and present the pertinent results of numerical simulations of the model. In chapter 3, I shall further discuss the choice of the inactivity mechanism for the VAEMG and then formalize the model's structure. This chapter will conclude with my developing a mean-field theory of the VAEMG. In chapter 4 I present the results of numerical simulations of the VAEMG model and compare these with the results of the theory developed in chapter 3. These results will focus on the two areas of the model indicated above: 1. how the performance and behaviour of an agent is affected when he has the ability to control his activity level; and 2. how the dynamics of the EMG model change when this modification is added. Finally, in chapter 5, I restate the pertinent results of the paper along with any possible implications they may have and then conclude by discussing the future work that the results suggest.

Chapter 2

The Evolutionary Minority Game

In this chapter I will review the Evolutionary Minority Game (EMG) model and discuss its behaviour with the use of numerical simulations. The work in this chapter is based upon [70] and will serve as a base for the rest of the thesis.

2.1 The basic EMG model structure

While the definition of the EMG given in chapter 1 might be enough to construct a numerical algorithm for the model, it is not adequate for any sort of an analytical treatment. In order to give a more precise description of the model and motivate a later analytical theory, a formal definition of the EMG is needed.

Let N be the (odd) number of agents in the game and let $i \in \{1, \ldots, N\}$ be the index that runs over this group. The game is played repeatedly for T rounds. On round $t \in \{1, \ldots, T\}$, each agent independently makes a binary decision, $\chi_{i,t} \in \{-1, +1\}$, about whether to buy (which will be arbitrarily labeled +1) or to sell (-1) the asset for that round. Once all agents have made their decisions the number who chose to buy and the number who chose to sell are calculated. The winning option, $\mu_t \in \{0, 1\}$, is defined as another binary variable whose value is 0 if $\sum_i \chi_{i,t} > 0$ (ie. sellers win) and +1 if $\sum_i \chi_{i,t} < 0$ (buyers win).

Now let the memory, M, be the number of past rounds that each agent can remember and $h_t^M = \mu_{t_M} \otimes \mu_{t-(M-1)} \otimes \ldots \otimes \mu_{t-2} \otimes \mu_{t-1}$ be the game's history of length M for round t. Given the binary nature of the game and a memory of length M, it follows that there are 2^M possible histories, h_t^M; each history (a binary string) will be labeled by its decimal equivalent, $j \in \{0, \ldots, 2^M - 1\}$.

An agent's trading decision, $\chi_{i,t}$, on round t is determined by the interaction between the current history, h_t^M, and his set of strategies. A strategy, $\mathcal{S}$, is a function that maps each of the possible 2^M histories onto some decision ($\mathcal{S} : \mu^M \to \delta$) that will be made given that history is the current history. Each strategy will be represented as a bit string of length 2^M where the j^{th} element, s_j, is the agent's trading decision given that the j^{th} history is the current history. For concreteness I have listed all of the possible histories for $M = 3$ along with a sample strategy in figure 2.1 below. If the history bit string labelled by $j = 3$ is encounterred during round t of the game (ie. $h_t^M = 011$) and an agent is playing this strategy, then that agent will choose to buy the asset ($s_3 = +1$) on that round.

Because of the binary nature of the decision process there are 2^{2^M} unique strategies for the 2^M possible histories; in the EMG agents only have access to 2 of these strategies. Let $\mathcal{S}^\dagger$ be the *trend* strategy. The j^{th} entry of this

j	h_t^3			s_j
0	0	0	0	-1
1	0	0	1	+1
2	0	1	0	+1
3	0	1	1	+1
4	1	0	0	-1
5	1	0	1	-1
6	1	1	0	+1
7	1	1	1	-1

Figure 2.1: The 8 possible histories, h_t^M, for $M = 3$ along with the 8 corresponding strategy elements, s_j, of a sample strategy, $\mathcal{S}$.

binary string holds the decision value, $\chi_{i,t}$, that would have put the agent in the minority room the last time the j^{th} history was encounterred in the game. Since the result of the game fluctuates between "buy" and "sell", $\mathcal{S}^{\dagger}$ is a dynamic strategy that will be modified after each round. While it might seem intuitive to play the trend strategy, if all agents took this view then they would all lose because of the minority rule. Heterogeniety amongst the agents is introduced by their ability to employ the *anti-trend* strategy, $\widetilde{\mathcal{S}^{\dagger}}$, which is the strategy anti-correlated with the trend strategy. Finally, associated with agent i is a *gene value*, $p_i \in [0, 1]$; this the probability that agent i will play the trend strategy.

At the beginning of the game each agent starts with a score of 0 and both the trend and anti-trend strategies are initialized. To seed the first round of the game an imaginary history, $h_1^M = \mu_{-M+1} \otimes \ldots \otimes \mu_0$, is created. Details of these initializations will be presented in the next section.

During each round t of the game agents select their trading decisions based upon their gene value. With probaility p_i, agent i chooses the option dictated by $\mathcal{S}^{\dagger}$ and with probability $1-p_i$ he chooses the other option. Once the minority decision is determined all those agents who chose the minority option have 1 point added to their score; those who chose the majority option have 1 point deducted from their score. At the end of every round the strategies are updated with the result of the round just played.

An agent will play the game until his score falls below some critical level, $Z_c < 0$, at which point he is "killed-off" and replaced with a new agent. This mechanism is the evolutionary feature of the EMG. If an agent's performance is poor then it is reasonable to assume that eventually he will lose all of his money and pull out of the market. The new agent is given an initial score of 0 and a new gene value which is a uniformly distributed random number with radius, $R \in [0, 1]$ (reflective boundary conditions are employed to ensure $p \in [0, 1]$) centerred around the p_i value for the killed-off agent. In this way a poorly performing gene value is replaced by one that is hopefully better able to compete within the complex group dynamic.

2.2 Numerical results

A numerical algorithm has been created and run over a wide range of parameter values simulating the behaviour of the EMG model. These simulations will serve two purposes: 1. to characterize the basic EMG system; and 2. by comparing the results with the original work on the EMG [70] I hope to confirm the accuracy of the algorithm so it can be used as a base upon which to construct the Variable Activity EMG algorithm.

Throughout this work the focus will primarily be on an agent's performance within the collective actions of the group. The crucial feature differentiating one agent from another is his gene value, p_i. In the EMG, p_i indicates agent i's willingness to follow trends; a $p_i \approx 1$ means that the agent is a trend follower, while $p_i \approx 0$ means the agent employs a contrarian trading strategy. By quantifying an agent's performance in terms of his gene value we will hopefully be able to shed light on what type of strategy is best suited for this type of supply-and-demand governed collective.

For any agent, i, the stochastic process, $\{R_{i,t}\}$, of his payoff is a supermartingale - ie. the EMG is a "less-than-fair" game. Because of the payoff function (+1 point per round for winning and -1 point for losing) and the minority rule structure of the game, the agent's expected payoff per round is negative and so it is expected that he will be killed-off within some finite time.

Instead of his point total, which will eventually fall to Z_c, I will measure an agent's performance by the number of rounds he plays before dying - ie. his lifetime. For this reason, the simulation results will focus on the distribution of average lifetimes, $L(p)$, for the group of agents.

To initialize the simulations, the current history, h_1^M, the trend, $\mathcal{S}^\dagger$, and anti-trend, $\widetilde{\mathcal{S}^\dagger}$ strategies, and the agents' gene values, p_i, were randomized from a uniform distribution of their appropriate realizations. Simulations were also performed using other initializations, but in all cases $L(p)$ was found to be invariant under these changes.

2.2.1 Time evolution of the EMG

In chapter 1 I discussed the intrinsic use of inductive reasoning in the EMG. Since this feature makes no assumption about the equilibrium of the system, it will be useful to investigate its dynamic properties and whether or not it *does* actually reach an equilibrium state.

I will describe the *state* of the system by $|n_d|$, the magnitude of the difference between the sizes of the two positions. In figure 2.2 is plotted two time seriess of $|n_d|$ for the EMG. We can see that the system exhibits a transient

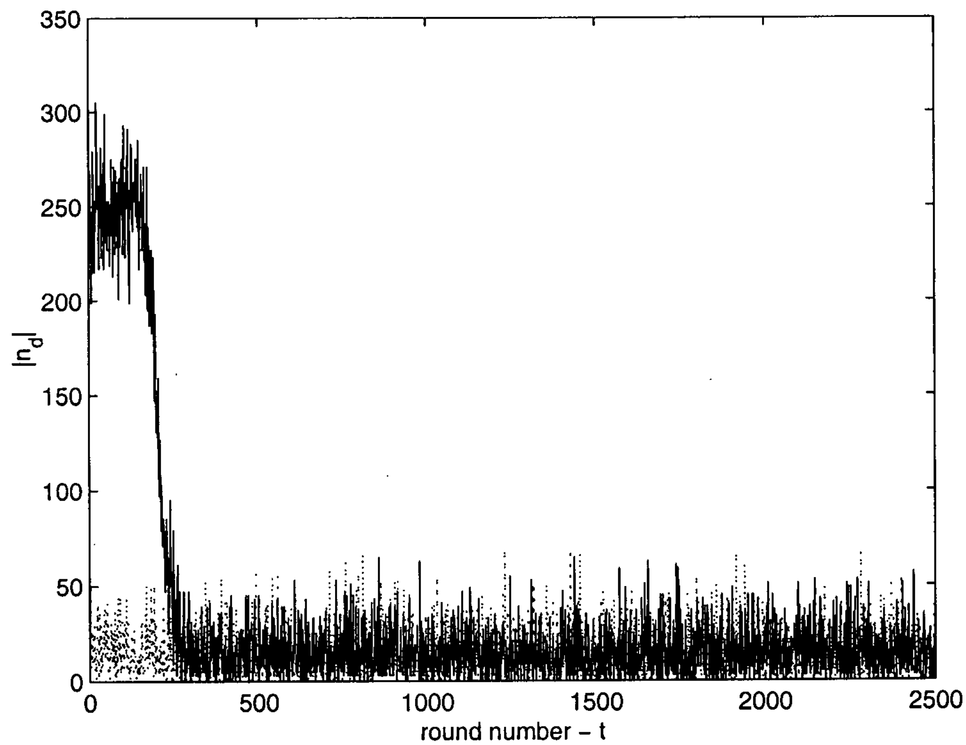

Figure 2.2: The time evolution of $|n_d|$ with $p_i = 0.25\ \forall i$ (—) and p_i randomly initialized ($\cdots$) for $N = 501$, $M = 3$, $Z_c = -100$, and $R = 1$.

response which rapidly decays away after approximately 200 rounds. The length and shape of this transient response appears relatively insensitive to the initial conditions of the system with one exception; when each agent's gene value is initially randomized the transient response is not clearly present. When the gene values are initially very similar to one another, the size of the majority group is very large (and therefore so is $|n_d|$); for this initial period most agents will lose a point on each round and so will die off in $\approx |Z_c|$ rounds. Once this first round of agents has been killed-off and replaced by agents with randomly distributed p_i values (since $R = 1$), $|n_d|$ then rapidly decays to a time-averaged value that remains constant for the rest of the game.

The state of the system in this "dynamic equilibrium" (in which the average of $|n_d|$ over some time interval $t \sim O(10^3)$ rounds) reveals a very interesting property of the model. If we examine $|\bar{n_d}|$, the time-averaged value of $|n_d|$, we will be able to determine how well the agents are utilizing the system's resources. As $|n_d| \to 1$ the size of the minority group becomes largest and the least wealth is lost from the system. In figure 2.3 I have plotted $|\bar{n_d}|$ for several values of N. To compare the efficiency of the EMG's equilibrium state with the "random" case, I have also plotted $|\bar{n_d}|$ for simulations in which the agents simply select their position at random (ie. independent of the game's history). On a log-log scale, both curves are very linear with exponents of 1/2 demonstrating that $|\bar{n_d}| \propto \sqrt{N}$ as we might expect.

What is interesting from figure 2.3 is the fact that the value of $|\bar{n_d}|$ for the EMG is significantly less than the "random" $|\bar{n_d}|$. This result tells us that the agents within the EMG naturally evolve (over the transient time of the system) into a state in which the systems' resources are used much

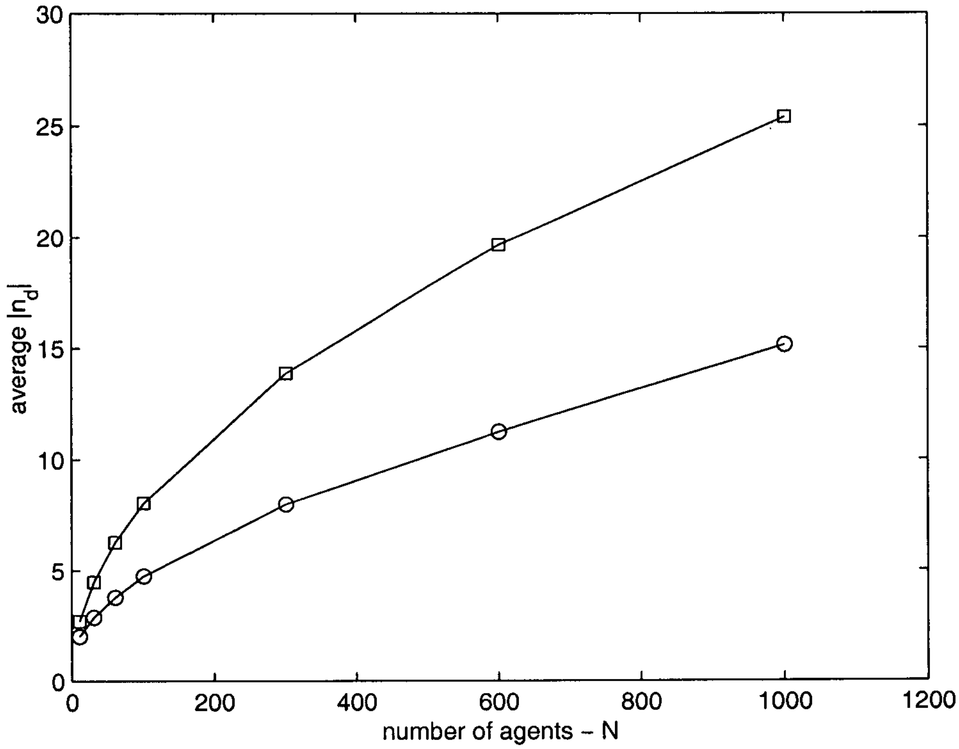

Figure 2.3: The equilibrium average size difference between the two positions for the EMG (○) and when the agents select their positions randomly (□). Other parameters for the simulations are $Z_c = -100$, $M = 3$, and $R = 1$.

more efficiently than if the agents simply made their decisions at random. This emergence of "co-operation" is most likely a result of the evolutionary feature of the model and fascinating due to the fact that, as will be mentioned below, the equilibrium properties of the model are independent the agent's memory, M, which one might expect to be a large source of the model's history dependence and therefore its emergent behaviour.

2.2.2 Equilibrium properties of $L(p)$

A sample $L(p)$ distribution is shown in Figure 2.4. The distribution is sym-

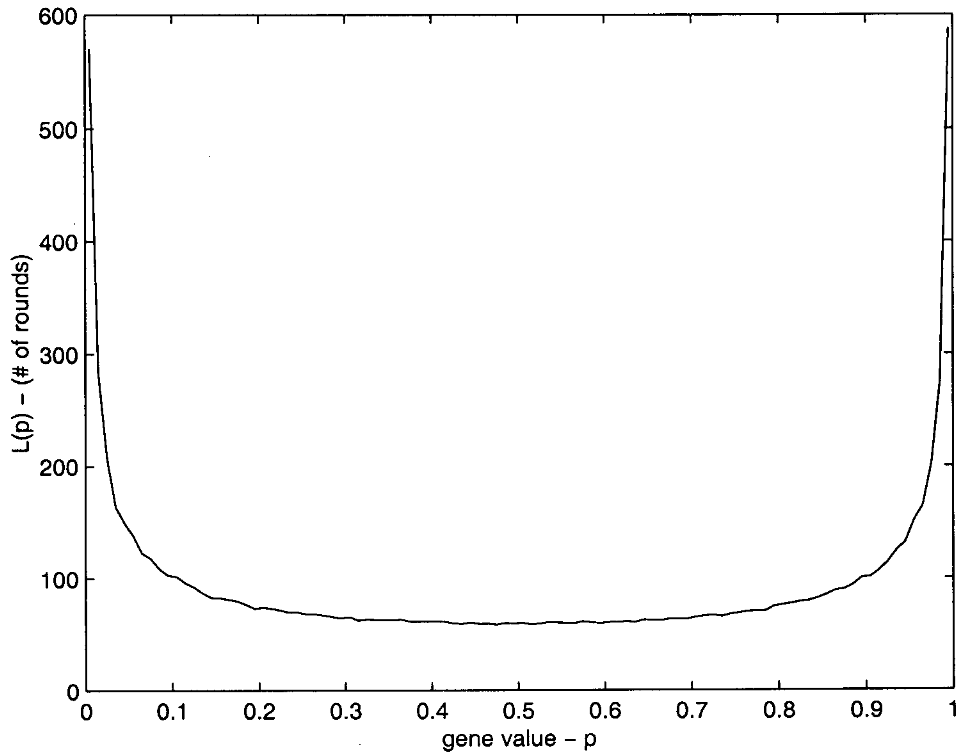

Figure 2.4: The distribution of average lifetimes, $L(p)$, for $N = 101$, $M = 3$, $T = 10^6$ rounds, $Z_c = -5$, and $R = 2$.

metric about $p_i = 1/2$ and is distinctly bimodal with sharp peaks at $p = 0$ and $p = 1$. Let p^* be an agent's optimal strategy - one that allows him to survive the longest. Interestingly, Figure 2.4 shows that this optimal strategy is to either always follow ($p^* = 1$) or go against ($p^* = 0$) the prevailing trend; if he uses any mixture of these strategies his performance will drop. In the case of $p = 1/2$ when an agent is most unsure of which strategy to choose, we can see that his performance is worst.

Figure 2.5 shows the dependence of $L(p)$ on the number of agents, N. As the number of agents increases the average lifetime of each agent increases also; this result can be understood if we consider an agent's expected payoff

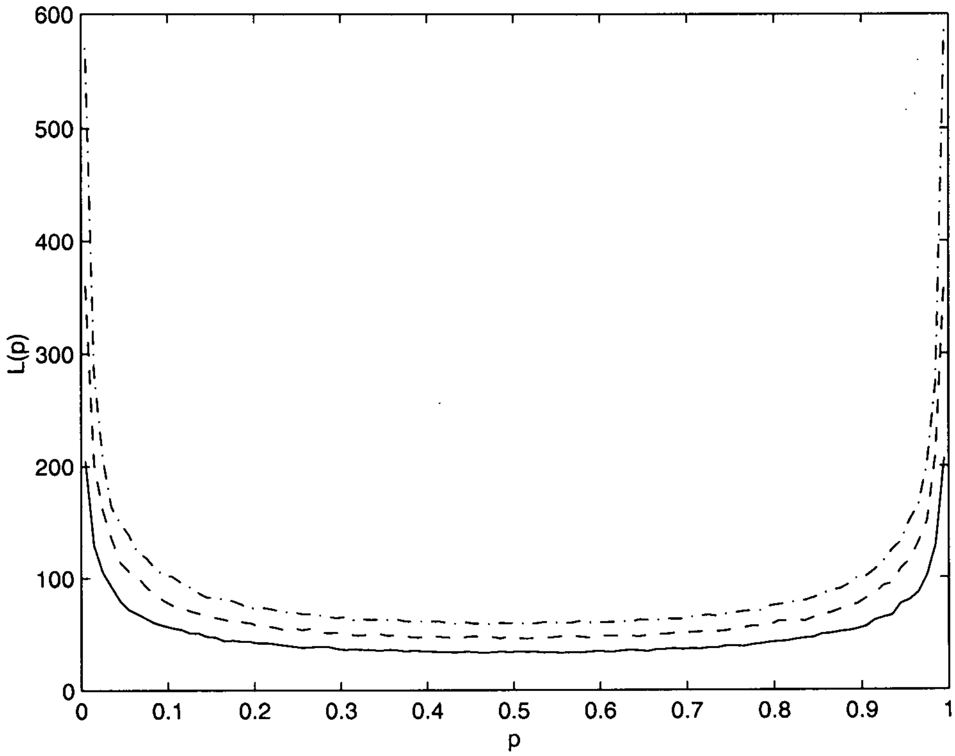

Figure 2.5: $L(p)$ for $M = 3$, $T = 10^6$, $Z_c = -5$, and $R = 2$ with N=31 (—); N=61 (---); and N=101 (-·-·-).

per round. From above we know that $|n_d|/N \propto 1/\sqrt{N}$ when the system is in equilibrium and so, as the number of agents increases, we should therefore expect the *fractional* difference between the size of the two options to go to zero. If this fractional difference decreases then so too will the expected payoff of the agents thereby increasing their average lifetimes.

Figure 2.6 shows the dependence of $L(p)$ on Z_c. As Z_c decreases it is intuitive that $L(p)$ should increase. Specifically we can see that $L(p)$ scales very well with the size of the cutoff score, $|Z_c|$, or $L(p) \sim |Z_c|$.

In addition, the dependence of $L(p)$ on M and R was checked. In all

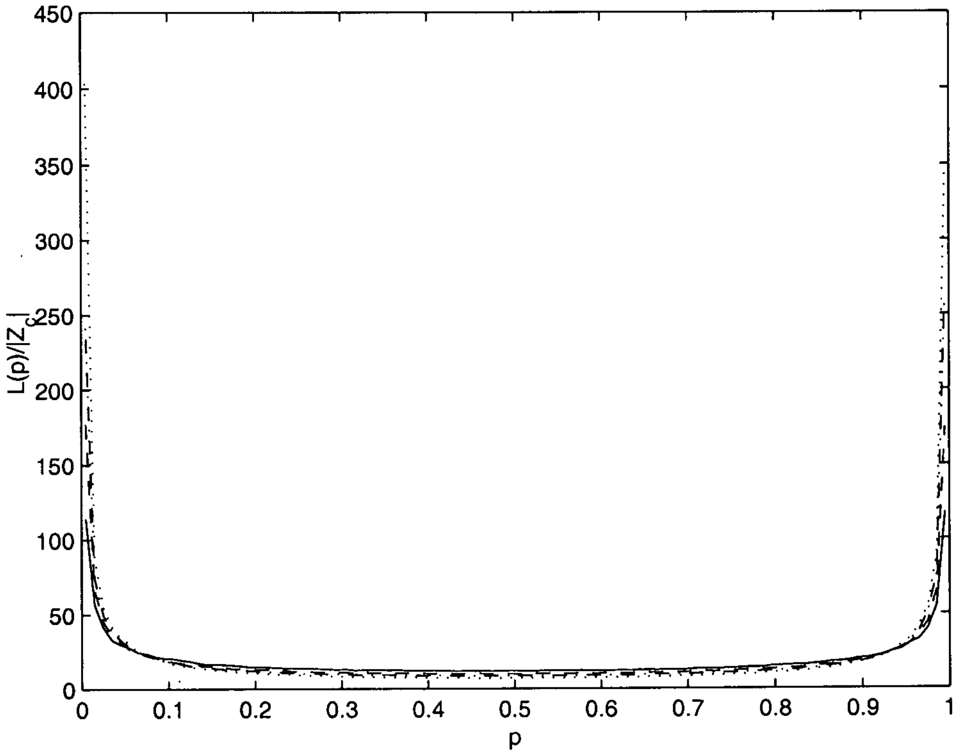

Figure 2.6: $L(p)/|Z_c|$ vs. p for Z_c = -5 (—), -10 (---), -25 (-·-·-), and -100 (···). $N = 101$, $M = 3$, $T = 10^6$, and $R = 2$.

cases there was no statistically significant dependence of $L(p)$ on either of these parameters.

Analysis of results

I will now give a simple explanation [73] for the cause of the bimodal and symmetric shape of $L(p)$ seen in figure 2.4. Consider the case $N = 3$ and assume that each agent can only take on the gene values, p =0, 1/2, or 1. I will label the agents 1, 2, and 3 and focus on the actions of agent 3.

Given the 3 possible values for p it follows that there are 9 possible gene value combinations, (p_1, p_2), for agents 1 and 2. Depending on the specific combination of (p_1, p_2) and his own gene value, p_3, agent 3 can expect a range of payoffs which are shown below in figure 2.7. The mean and variance for

(p_1, p_2)	(0,0)	(0,$\frac{1}{2}$)	(0,1)	($\frac{1}{2}$,0)	($\frac{1}{2}$,$\frac{1}{2}$)	($\frac{1}{2}$,1)	(1,0)	(1,$\frac{1}{2}$)	(1,1)
$p_3 = 0$	-1	-1	-1	-1	-$\frac{1}{2}$	0	-1	0	+1
$p_3 = \frac{1}{2}$	0	-$\frac{1}{2}$	-1	-$\frac{1}{2}$	-$\frac{1}{2}$	-$\frac{1}{2}$	-1	-$\frac{1}{2}$	0
$p_3 = 1$	+1	0	-1	0	-$\frac{1}{2}$	-1	-1	-1	-1

Figure 2.7: All possible expected payoffs for agent 3 when $N = 3$ given that each agent can only access the states p =0, 1/2, and 1.

each of the rows of figure 2.7 are shown in figure 2.8.

p_3	mean payoff	variance
0	-$\frac{1}{2}$	$\frac{2}{9}$
$\frac{1}{2}$	-$\frac{1}{2}$	$\frac{1}{9}$
1	-$\frac{1}{2}$	$\frac{2}{9}$

Figure 2.8: The mean payoff and corresponding variance for each of agent 3's gene values when $N = 3$ and each agent only has access to the states p =0, 1/2, and 1.

From figure 2.8 we can see, as expected, that the mean payoff for agent 3 is the same for each of his gene values (ie. there is no *a priori* best gene value), but that the variance is greater for $p_3 = 0$ and $p_3 = 1$ than for $p_3 = 1/2$. Eventhough there are more situations in which agent 3 is *guaranteed* to choose the majority position (and therefore be penalized 1 point) with either p_3 =0 or 1, the larger variances for these gene values indicate that the agent also has a greater probability of choosing the minority position by out-guessing the other two agents; he is therefore more likely to live for a greater number of rounds before he is killed-off.

2.3 Discussion of results

The results of simulations of the EMG presented in this chapter have demonstrated some of the interesting properties of the model. First, it has been shown that agents within the EMG framework naturally organize themselves into a state in which the systems' resources are utilized more efficiently than if the agents simply chose their positions at random (ie. by flipping a coin). This self-organized "co-operation" was shown to be independent of the system's initial state and, since its equilibrium state was found to be independent of the agent's memory, is most likely a result of the history-dependent evolutionary mechanism of the model.

Also significant is the bimodal and symmetric nature of the distribution of average lifetimes, $L(p)$. Within the type of competitive group dynamic formed by the EMG, it has been shown that an agent's optimal strategy is to either always or never follow the trend set by the group. This result has been explained using the simple arguement that the variance for an agent's payoff is a minimum for $p = 1/2$ and a maximum for $p = 0$ and $p = 1$; when an agent is in one of these extreme strategies he therefore has the greatest probability of choosing the minority position and surviving the longest.

While the unrealistic features of the MG have been discussed in chapter 1, there is one point specific to the EMG that I feel should be emphasized further. It has already been mentioned several times that the EMG is a super-Martingale and so the expected payoff for each agent is always negative. If the results of the model are discussed in financial terms then one immediately encounters the problem *If, on average, an agent expects to lose money, regardless* *of which position he takes, then why would he trade in the first place?*

In reality, the EMG is actually more a realistic model of a group of agents either betting "red" or "black" on a roulette wheel than a group of agents trading a financial asset in a market. The only agents that would play the EMG are those that are risk-preferring - ie. agents would only participate for the thrill of the game and not with the realistic expectation of increasing their wealth. Since the lifetime is the only decent measure of such a risk-preferring agent, the bimodal shape of the distribution of average lifetimes is not unexpected; the strategies $p = 0$ and $p = 1$ are the riskiest strategies and so it is reasonable that they are also the best performing.

In keeping with the theme of developing the EMG into a more realistic market model, an interesting (and essential) extension would be to incorporate a payoff function in which the expected payoff is positive and the agents therefore have an incentive to participate. By defining a "wealth" as an agent's accumulated score and incorporating an evolutionary mechanism which replaces the gene values of a certain percentage of the worst performing agents every $z \gg 1$ rounds of the game, one can then examine the group's performance with the final distribution of wealth and of gene values.

Because of the sub-Martingale property of this proposed model it is likely that using either $p = 0$ or $p = 1$ will no longer be optimal - ie. they will not maximize an agent's wealth. In fact, it is quite possible that $p = 1/2$ will emerge as an agent's optimal gene value due to its lesser probability of creating a situation in which the agent is guaranteed to choose the minority room and lose (see figure 2.7).

Even though the current binary payoff function is financially unrealistic, I will use it in this work and leave the above modifications for a later study.

Chapter 3

The Variable Activity Evolutionary Minority Game

Results of simulations of the EMG from chapter 2 have shown that the model is well-suited for demonstrating properties of an individual agent within the Minority Game framework. Specifically, it has been shown that an agent's optimal strategy (defined as the gene value which maximizes his expected lifetime) is to always play the same *pure* strategy - either always follow the previous winning option ($p = 1$), or always go against this trend ($p = 0$).

Several questions arise if the agents are given the ability to abstain from playing. Firstly, will the symmetry of the EMG be broken? It has been shown in chapter 2 that the trend-following and contrarian pure strategies of the EMG are equally optimal. It is interesting to ask whether this situation will persist, or if the ability to be inactive will cause one of these strategies to become preferable over the other. Secondly, what will the optimal strategy be when inactivity is introduced; will this strategy be "pure" such as *always play* or *never play*, or will it be some mixture of the two; are there symmetries of agent performance over the strategy space; and are there conditions under which an agent's optimal strategy can change?

These are some of the questions that have motivated the development of the Variable Activity Evolutionary Minority Game (VAEMG). In this chapter I will discuss and formalize the structure of the VAEMG model and specifically the structure of the mechanism used to generate agent inactivity within the standard EMG framework. With the VAEMG I have attempted to model this mechanism so as to capture some realistic features of a group of competitive and adaptive agents while keeping the model as simple as possible. The chapter will conclude using this formal definition to develop an analytical theory of the VAEMG which will later be used as a comparison with the results of numerical simulations of the model.

3.1 Choice of activity structure

The first problem that must be addressed is the way in which to model agent inactivity within the framework of the EMG. The possibilities for this choice are endless. For instance, we could use:

1. a very simple "group" condition where specific agents are randomly selected to sit-out on each round,

2. inactivity spreading via a bond percolation mechanism like in the Cont-Bouchaud model [31], or

3. some more realistic mechanism of controlling inactivity at the single

agent level through feedback of the agent's past performance

For this study it would be ideal if the inactivity mechanism generated emergent properties of the system, but also was very simple. With choice (1) above, while the global inactivity structure is very simple, it would unlikely produce any interesting results since it lacks any mechanism for generating correlations between the individual agents. With choices (2) and (3), mechanisms exist to cause the necessary correlations, but in both cases they are overly complex and will likely produce results that are difficult to analyze.

The EMG model provides an ideal starting point for this goal. As has already been noted, the evolutionary nature of the gene value, p in the standard EMG allows us to examine how an agent's performance depends on his strategy selection. Even though its structure is relatively simple, this evolution is important due to the feedback that introduces into the system. It is hoped that, if the agent's activity level is modeled in an analogous way, the modified model will produce interesting emergent behaviour (such as a relationship between agent performance and activity level) while remaining relatively easy to treat analytically.

3.2 Model structure

The play of the VAEMG is the same as the EMG with several exceptions. In the VAEMG, each agent i ($\forall i \in \{1, \ldots, N\}$) has an activity parameter, $\alpha_i \in [0, 1]$, which is simply that agent's probability of being active during a round. In an analogous way with p_i in the standard EMG, α_i is an evolutionary variable and therefore introduces feedback (memory dependence) into an agent's decision on his activity status.

The game proceeds in two independent steps. In the first step, each agent "chooses" whether or not to be active based upon his activity parameter, α_i. Once these decisions are made, the total number of active agents, $N_a \leq N$, is determined. While not strictly necessary, we have chosen to exclude the possibility of having an even number of active agents (ie. $N_a \in \{1, 3, \ldots, N\}$). In addition to keeping the model structure as close to the original EMG as possible and preventing us from having to deal with the case when there is no distinct minority group ($n_d = 0$), the exclusion of N_a-even, as will be shown in a later section, significantly reduces the complexity of the analytical theory. In the numerical algorithm the N_a-even condition is enforced simply by requiring the agents to replay a round's first step until an odd number of active agents is obtained.

In the second step of the round the active agents play a normal round of the EMG with $N = N_a$. Each of the agents chooses either to "buy" or "sell" based upon the past history and their individual strategy preference, p_i, and then is either rewarded or penalized based upon whether their choice puts them in the minority or majority group, respectively.

As the VAEMG is currently structured, there is no incentive to participate in the second step of the game. If an agent decides to be inactive for the round, then he simply sits out and waits for the next round. While he does have the possibility of choosing the minority option and raising his score, his expected payoff, given that he plays, is negative; he would therefore be better off by choosing never to play.

To prevent this "trivial" solution (all agents always inactive) I will introduce a global parameter, $I \in [-1, 0]$, to act as a penalty to an agent's score for being inactive during a round. More than just a mathematical convenience, I can be thought of as a parameter that models the effect of inflation on the system. To explain this point, imagine a group of agents within a financial market who own cash and a single type of asset. When an agent is confident (in his own mind) about the future movement of the asset he will "play" the market and either buy or sell some quantity of it and in turn change the quantity of his cash reserve. If, on the other hand, he is unsure of the asset's future movements he can choose to avoid the risk of playing the market and hold the quantities of his cash and asset constant. While this second option is less risky than the first, it will not entirely preserve his wealth. Because of inflation, the value of his cash, which could have been invested in the form of the asset, will be worth less in the future and so his net worth will effectively decrease.

The game proceeds until agent i's score falls *strictly* below the cut-off score, Z_c. As in the EMG the agent is then killed-off and replaced by a new agent with an initial score of zero and redistributed values of both α_i *and* p_i. These redistributions are uncorrelated random variables uniformly distributed around the old values of α_i and p_i with radii $R_\alpha \in [0, 1]$ and $R_p \in [0, 1]$, respectively.

3.3 Theory of the VAEMG

Like the standard EMG, the VAEMG is a super-martingale for all values of $I < 0$. Since each agent will eventually die (for $|Z_c| < \infty$) that agent's performance is best measured by his lifetime. In the VAEMG, however, there are two evolutionary variables, p_i and α_i, and since correlations between these variables may emerge, the study will focus on the generalized average lifetime distribution, $L(\alpha_i, p_i)$.

The goal of this section is to derive an analytical expression for $L(\alpha, p)$ (the subscript i has been omitted since the theory should be general for *any* agent within the group). The derivation will mirror that of Lo et al. [71] and will use a mean-field formulation. Because agents within the VAEMG only interact through its *global* history, the system should be ammenable to a such a mean-field treatment.

3.3.1 Derivation of an expression for the average lifetime distribution, $L(\alpha, p)$

I will begin by defining the expected payoff to an agent in states α and p as $\langle \Pi_{\alpha p} \rangle$. Since $\langle \Pi_{\alpha p} \rangle \leq 0$ it is expected that approximately $1/|\langle \Pi_{\alpha p} \rangle|$ rounds will be required for the agent to lose 1 point. Because the agent is killed-off when his score, S, reaches $Z_c < 0$ points, his average lifetime, $L(\alpha, p)$, should be given by

$$L(\alpha, p) = \frac{Z_c}{\langle \Pi_{\alpha p} \rangle}$$

The above equation is actually only valid for the standard EMG where an agent's score can only change by an integer amount and therefore $\langle S_c(\alpha,p)\rangle = Z_c$ is a constant (where $\langle S_c(\alpha,p)\rangle$ is, on average, the agent's score when he is killed-off). In the VAEMG, if an agent is inactive during a round then his score will change by a non-integer amount, $I \in [-1,0]$ and therefore $\langle S_c(\alpha,p)\rangle$ will most likely not equal Z_c. The average number of rounds required to reach the cut-off score (ie. the average lifetime) will therefore be slightly modified in the VAEMG to

$$L(\alpha,p) = \frac{\langle S_c(\alpha,p)\rangle}{\langle \Pi_{\alpha p}\rangle} \tag{3.1}$$

During each round of the VAEMG, there are 3 possible moves for an agent in states α and p:

1. He can choose to be inactive. This choice occurs with probability $1-\alpha$ and results in a payoff of I points.

2. He can choose to participate in the round and then choose the minority option. This occurs with probability $\alpha P_w(\alpha,p)$ (where $P_w(\alpha,p)$ is the probability that an agent with activity level α and gene value p chooses the winning option *given that he is active*) and results in a payoff of +1.

3. He can choose to play and then choose the majority option. This occurs with probability $\alpha(1-P_w(\alpha,p))$ and results in a payoff of -1 point.

The expected payoff to the agent is therefore

$$\langle \Pi_{\alpha p}\rangle = \alpha(2P_w(\alpha,p)-1) + I(1-\alpha) \tag{3.2}$$

Analytical form for the expected cutoff score, $\langle S_c(\alpha, p)\rangle$

To derive an expression for $\langle S_c(\alpha, p)\rangle$ I will examine the actions of only one agent. It will be assumed that this agent never chooses the minority option (ie. the agent is always inactive or active but chooses the majority position) thereby neglecting the effect of positive score changes. This assumption should cause no loss of generality since positive changes are always integer valued and therefore do not contribute to the variability in $\langle S_c(\alpha, p)\rangle$.

The expected payoff of this agent is therefore $\alpha P_w(\alpha, p) + (1-\alpha)I$ and so, on average, the number of rounds it will take for his score to drop to Z_c will be $Z_c/(\alpha P_w(\alpha, p) + (1-\alpha)I)$. Since the agent will not be killed-off until $S < Z_c$ this number of rounds must be rounded-up to the nearest integer (a fractional number of rounds is not allowed). The expectation of the agent's score, $\langle S_c(\alpha, p)\rangle$, at this point will simply be this number of rounds multiplied by his expected payoff, or

$$\langle S_c(\alpha, p)\rangle = (\alpha P_w(\alpha, p) + (1-\alpha)I) \times ceil\left[\frac{Z_c}{\alpha P_w(\alpha, p) + (1-\alpha)I}\right] \qquad (3.3)$$

where $ceil(x)$ is defined as the smallest integer no smaller than x.

3.3.2 Derivation of an expression for $P_w(\alpha, p)$

The derivation of $L(\alpha, p)$ is now reduced to finding an expression for $P_w(\alpha, p)$. Towards this goal, I will focus on some round, $t \gg 1$, when the system has reached its equilibrium state. At this point I will consider the second step of the round in which there are N_a active agents and, specifically, the actions of the i^{th} agent in a pool of $N_a - 1$ background agents.

I define $F_{N_a}(n)$ as the probability that n of the total N_a active agents choose the option (buy or sell) predicted by the *trend* strategy. Also let $G^i_{N_a-1}(n)$ be the probability that n agents, chosen from the background pool, choose the option predicted by the trend strategy. $F_{N_a}(n)$ and $G^i_{N_a-1}(n)$ are then related by

$$F_{N_a}(n) = p_i G^i_{N_a-1}(n-1) + (1-p_i) G^i_{N_a-1}(n) \tag{3.4}$$

which simply states that there are two ways in which n agents choose the trend strategey; either if agent i chooses this strategy along with $n-1$ of the background agents, or if n of the background agents choose the strategy while agent i chooses the anti-trend strategy.

Agent i will choose the winning option if one of the two following situations occur:

1. Either agent i chooses the trend strategy and no more than $(N_a - 3)/2$ of the background agents choose this same option, or

2. Agent i chooses the anti-trend strategy while at least $(N_a + 1)/2$ of the background agents choose the trend strategy

The probability that agent i chooses the minority (winning) option, $p_w(N_a)$, given that he is both active *and* there are N_a active agents is therefore

$$p_w(N_a) = p_i \sum_{n=0}^{\frac{N_a-3}{2}} G^i_{N_a-1}(n) + (1-p_i) \sum_{n=\frac{N_a+1}{2}}^{N_a-1} G^i_{N_a-1}(n) \tag{3.5}$$

Equation (3.5) is i-dependent. To remove this dependence and generalize the result for all agents, (3.4) must be rewritten as

$$F_{N_a}(n) = G^i_{N_a-1}(n) - p_i \left(G^i_{N_a-1}(n) - G^i_{N_a-1}(n-1) \right) \tag{3.6}$$

and note its boundary conditions

$$F_{N_a}(0) = (1 - p_i)G^i_{N_a-1}(0) \tag{3.7}$$

$$F_{N_a}(N_a) = p_i G^i_{N_a-1}(N_a - 1) \tag{3.8}$$

Now summing over both sides of (3.6) from $n = 1$ to $n = (N_a - 3)/2$ and cancelling terms gives

$$\sum_{n=1}^{(N_a-3)/2} F_{N_a}(n) = \sum_{n=1}^{(N_a-3)/2} G^i_{N_a-1}(n) + p_i G^i_{N_a-1}(0) - p_i G^i_{N_a-1}(\frac{N_a - 3}{2})$$

Using the b.c. (3.7) the above equation reduces to

$$\sum_{n=0}^{(N_a-3)/2} G^i_{N_a-1}(n) = \sum_{n=0}^{(N_a-3)/2} F_{N_a}(n) + p_i G^i_{N_a-1}(\frac{N_a - 3}{2}) \tag{3.9}$$

In an analagous manner, (3.6) can be summed from $n = (N_a + 1)/2$ to $n = N$. After the appropriate cancellations and use of b.c. (3.8) we are left with

$$\sum_{n=(N_a+1)/2}^{N_a-1} G^i_{N_a-1}(n) = \sum_{n=(N_a+1)/2}^{N_a-1} F_{N_a}(n) - p_i G^i_{N_a-1}(\frac{N_a - 1}{2}) \tag{3.10}$$

Substituting (3.9) and (3.10) into (3.5) and using the relation (from (3.4))

$$p_i G^i_{N_a-1}((N_a - 3)/2) = F_{N_a}((N_a - 1)/2) - (1 - p_i)G^i_{N_a-1}((N_a - 1)/2)$$

gives

$$p_w(N_a) = p_i \sum_{n=0}^{\frac{N_a-1}{2}} F_{N_a}(n) + (1 - p_i) \sum_{n=\frac{N_a+1}{2}}^{N_a} F_{N_a}(n) - 2p_i(1 - p_i)G^i_{N_a-1}\left(\frac{N_a - 1}{2}\right) \tag{3.11}$$

To remove all i-dependence from $p_w(N_a)$ the third term on the RHS of (3.11) needs to be expressed in terms of $F_{N_a}(n)$. Again, rewriting (3.4) we have

$$p_i G^i_{N_a-1}(n-1) = F_{N_a}(n) - (1-p_i) G^i_{N_a-1}(n) \tag{3.12}$$

Applying (3.12) to itself then results in the relation

$$p_i G^i_{N_a-1}(n-1) = F_{N_a}(n) - \frac{1-p}{p}\left[F_{N_a}(n+1) - (1-p_i) G^i_{N_a-1}(n+1)\right]$$

and after successive applications, we are left with

$$\begin{aligned} p_i G^i_{N_a-1}(n-1) = F_{N_a}(n) - \frac{1-p}{p} F_{N_a}(n+1) + \left[\frac{1-p}{p}\right]^2 F_{N_a}(n+2) + \ldots \\ + \left[\frac{p-1}{p}\right]^{N_a-n} F_{N_a}(N_a) \end{aligned}$$

or simply

$$G^i_{N_a-1}(n-1) = \frac{1}{p} \sum_{j=0}^{N_a-n} \left[\frac{p-1}{p}\right]^j F_{N_a}(n+j) \tag{3.13}$$

Because of the factor $[(p-1)/p]^j$, equation (3.13) is not numerically practical for $p \approx 0$. To correct this problem, we write (3.4) as

$$(1-p_i) G^i_{N_a-1}(n) = F_{N_a}(n) - p_i G^i_{N_a-1}(n-1) \tag{3.14}$$

After recursively applying Eq. (3.14) to itself as was done with Eq. (3.12) we get

$$G^i_{N_a-1}(n) = \frac{1}{1-p} \sum_{j=0}^{n} \left[\frac{p}{p-1}\right]^j F_{N_a}(n-j) \tag{3.15}$$

which is much more usable in the region $p \approx 0$.

Together, equations (3.11), (3.13), and (3.15) give the probability, $p_w(N_a)$, that any agent chooses the minority option given that he is both active *and* there are N_a active agents during the round in question. But $N_a \in \{1, 3, \ldots, N\}$ is variable. To find the *total* probability that the agent chooses the minority option, it is necessary to perform a weighted sum over all possible states, N_a.

Towards this goal I will now focus on the first step of the round and, specifically, on the actions of the i^{th} agent within a pool of $N-1$ background agents (all N agents are still "active" at this stage). Let $\Theta_N(N_a)$ be the probability that ther are N_a active agents, out of a possible N. Also let $\Phi^i_{N-1}(N_a - 1)$ be the probability that, excluding the i^{th} agent, there are $N_a - 1$ active background agents during the round. Simply by definition, the probability that agent i selects the minority option given that he is active is then

$$P_w(\alpha, p) = \sum_{N_a=3}^{N} \Phi^i_{N-1}(N_a - 1) p_w(N_a) \tag{3.16}$$

In a similar way to Eq (3.4), $\Theta_N(N_a)$ and $\Phi^i_{N-1}(N_a - 1)$ are related by

$$\Theta_N(N_a) = \alpha_i \Phi^i_{N-1}(N_a - 1) + (1 - \alpha_i)\Phi^i_{N-1}(N_a) \tag{3.17}$$

which states that there can be N_a active agents if agent i is active along with $N_a - 1$ of the background agents, or if agent i is inactive while N_a background agents *are* active. The expression for $P_w(\alpha, p)$ is conditional on agent i being active; if he is inactive then $P_w(\alpha, p) = 0$ by definition. The 2^{nd} term on the RHS of (3.17) assumes that agent i is inactive and it therefore makes no contribution to $P_w(\alpha, p)$. For our purposes then, (3.17) reduces to

$$\Phi^i_{N-1}(N_a - 1) = \frac{1}{\alpha}\Theta_N(N_a) \tag{3.18}$$

and so the final, i-independent form of $P_w(\alpha, p)$ is

$$P_w(\alpha, p) = \frac{1}{\alpha} \sum_{N_a=3}^{N} \Theta_N(N_a) p_w(N_a) \tag{3.19}$$

3.3.3 Analytical forms for $\Theta_N(N_a)$ and $F_{N_a}(n)$

The final question that must be answerred is how to deal with $\Theta_N(N_a)$ and $F_{N_a}(n)$. Qualitatively, $\Theta_N(N_a)$ and $F_{N_a}(n)$ can be regarded as the sum of N and N_a independent actions, respectively. At this time I will restrict $N \gg 1$ and assume that $N_a \gg 1$ also. Since $\alpha_i, p_i \in [0, 1]$ the Central Limit Theorem tells us that $\Theta_N(N_a)$ and $F_{N_a}(n)$ should both be approximately gaussian shaped with means, μ_α and μ_p, and variances, σ_α^2 and σ_p^2, respectively, given by

$$\mu_\alpha = N \int_\alpha \int_p \alpha P(\alpha, p) dp d\alpha \qquad \sigma_\alpha^2 = N \int_\alpha \int_p \alpha(1-\alpha) P(\alpha, p) dp d\alpha \tag{3.20}$$

$$\begin{aligned} \mu_p &= \sum_{N_a} \Theta_N(N_a) N_a \int_\alpha \int_p p P(\alpha, p) dp d\alpha \\ \sigma_p^2 &= \sum_{N_a} \Theta_N(N_a) N_a \int_\alpha \int_p p(1-p) P(\alpha, p) dp d\alpha \end{aligned} \tag{3.21}$$

where $P(\alpha, p)$ is simply the joint frequency distribution of α and p and is given by

$$P(\alpha, p) = \frac{L(\alpha, p)}{\int_\alpha \int_p L(\alpha, p) d\alpha dp} \tag{3.22}$$

when the system is in equilibrium (this will be experimentally verified in a later section).

3.3.4 Solving for the analytical form of $L(\alpha, p)$

The derivation is now complete. Together, (3.1-3.3), (3.11), (3.13), (3.15), and (3.19-3.22) form a self-consistent set of equations for $L(\alpha, p)$. Unfortunately this system is implicit and so $L(\alpha, p)$ cannot be solved directly. To solve for $L(\alpha, p)$ I will use the following iterative routine:

1. Assume an initial form for $L(\alpha, p)$.

2. Assume an initial gaussian form (with $\mu = \sigma^2 = N/2$) for $\Theta_N(N_a)$ for use in Eq. (3.22).

3. Calculate μ_α, μ_p, σ_α^2, and σ_p^2 using equations (3.20) and (3.21).

4. Using equations (3.13) and (3.15), calculate $p_w(N_a)$ from (3.11)

5. Using (3.19) calculate $P_w(\alpha, p)$

6. Using (3.3) calculate the cutoff score, $\langle S_c(\alpha, p)\rangle$

7. Calculate the expected payoff, $\langle \Pi_{\alpha p}\rangle$, (3.2) and then find a corrected form for $L(\alpha, p)$ using (3.1)

8. Using the new $L(\alpha, p)$, repeat steps 3-6 until the difference between successive $L(\alpha, p)$'s falls within some convergence criterion.

Chapter 4

Results and Discussion

In this chapter I will examine the properties and behaviour of the VAEMG model developed in chapter 3. As with the EMG model of chapter 2, I will focus on results of numerical simulations of the model and then supplement and compare these with the theory developed in section 3.3. The results will focus on two aspects of the model: 1. the performance of an agent when he is able to control his activity level; and 2. how the properties of the EMG model react to a variable level of agent activity.

In section 4.1 I will begin the chapter with a short discussion about the details of initializing the routines for generating both the numerical and theoretical results. In section 4.2 I will discuss the transient response of the VAEMG model. In section 4.3 I will then present and discuss results concerned with the activity-related properties of a single agent. In order to focus on this aspect of the model, I will use a version of the VAEMG in which the gene value, p, is non-evolutionary for all agents. As will be explained later, this formulation reduces the complexity of the model and allows us to focus more easily on the activity dependence of an agent's behaviour. Finally, in section 4.4 I will use the full version of the VAEMG model developed in section 3.2 to examine how the addition of a variable activity level affects global properties of the EMG model.

4.1 Setup of numerical routines

The data presented in this chapter have been generated from numerical routines. The purpose here is to simply describe some aspects of the setup and any relevant initial conditions of these routines.

4.1.1 Experimental

The "experimental" data (data from numerical simulations of the VAEMG model) have been produced using a routine built-up from that used to simulate the EMG in chapter 2. As with those simulations, the initial history bit string, h_1^M, the trend and anti-trend strategies ($\mathcal{S}^\dagger$ and $\widetilde{\mathcal{S}^\dagger}$ respectively), and an agent's activity level, $\alpha_i \in [0,1]$ were all randomly initialized. In addition, an agent's gene value, $p_i \in [0,1]$, was randomly initialized in section 4.4 where the full VAEMG model is simulated. In section 4.3 however, p_i is non-evolutionary and so it has been set the same for all agents; the choice of this constant will be discussed further in section 4.3.

Simulations of the VAEMG are computationally very large. Whereas the number of rounds, T, used in a simulation of the EMG in chapter 2 was 10^6, I will show in section 4.3 that T must be much larger (typically $T \sim O(10^8 -$

10^9)) for the VAEMG model. As a result of this increased computational size, I have used the $ran2()$ random number generator found in *Numerical Recipes in C* [74]. This random number generator is of the L'Ecuyer type with a Bays-Durham shuffle and, most importantly, has a very long period ($> 2 \times 10^{18}$) which should prevent any serial correlations between rounds.

4.1.2 Theoretical

The routine, given in section 3.3.4, for generating the theoretical distribution of average lifetimes has several points that must be discussed. First, the "continuous" variables, α and p (where appropriate), have each been discretized into 100 evenly spaced points over the interval (0,1). Second, an initial average lifetime distribution was required to seed the numerical routine; for simplicity reasons a constant, normalized distribution was used. The sensitivity of the algorithm to this initial distribution was tested using various other forms and, in all cases it was found that the final distribution was independent of this choice.

The precision of the final distribution is dependent on the convergence criterion used in the algorithm. Let $L_i(\alpha, p)$ be the i^{th} approximation to the actual theoretical distribution. We define

$$\gamma_i = \sqrt{\sum_{\alpha p} \left(L_i(\alpha, p) - L_{i-1}(\alpha, p)\right)^2}$$

to be a measure of the difference between the $i-1^{th}$ and i^{th} iterations of the routine (and the $\sum_{\alpha p}$ is a sum over all possible, discrete combinations of α and p). When the convergence criterion is set to be $\gamma_i \leq 10^{-6}$ there is no visible change in the distribution between steps $i-1$ and i. For safety sake I will use $\gamma_i \leq 10^{-10}$ as the convergence criterion throughout this study.

4.2 Transient response of the VAEMG

While I will mainly focus on the equilibrium properties of the VAEMG model, I would also like to characterize its approach to this state. With the EMG the approach to the dynamic equilibrium was quantified using $|n_d|$, the magnitude of the difference between the number of agents who selected "buy" and who selected "sell". With the VAEMG there are three possible outcomes each round (buy, sell, and hold) and so two quantities are needed: the magnitude of the difference between the number of buyers and sellers, $|n_d|$; and the magnitude of the difference between the number of active and inactive agents, $|2N_a - N|$. Time series for each of these quantities are shown in figure 4.1. The transient responses of both $|n_d|$ and $|2N_a - N|$ appear to have died out completely by $t \approx 1000$ rounds.

There are two interesting points from figure 4.1. First, from figure 2.2 we know that, within the EMG framework, $|n_d|$ will not display a rapidly decaying transient if p_i is randomly initialized; figure 4.1 re-confirms this result for the VAEMG. What is more interesting, however, is the fact that $|2N_a - N|$ *does* rapidly evolve out of this randomly initialized state with a characteristic time $\sim O(|Z_c|)$. As will hopefully become obvious later, this spontaneous evolution of $|2N_a - N|$ out of a state in which α_i is randomly distributed (in comparison to the evolution of $|n_d|$ *into* a state which behaves similarly to one

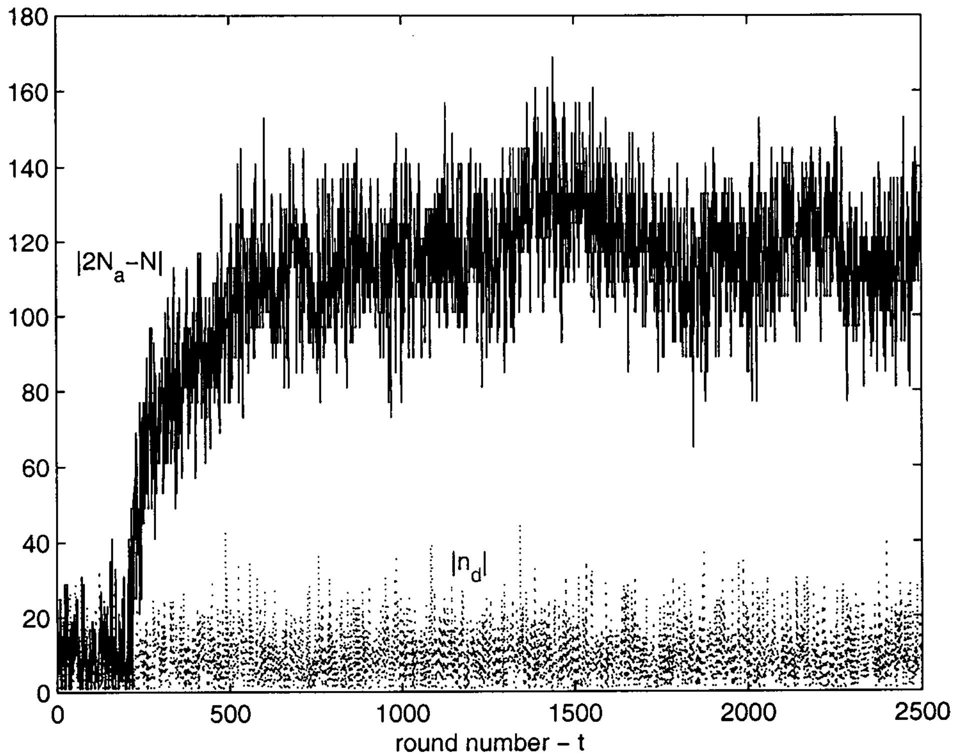

Figure 4.1: Time series for $|n_d|$ and $|2N_a - N|$ for $N = 301$, $M = 3$, and $I = -0.5$ with both p_i and α_i evolutionary and randomly initialized.

in which p_i is randomly distributed) is a result of an agent's performance being asymmetrically dependent on α_i. Second, once $|2N_a - N|$ has evolved into its dynamic equilibrium level, it appears to undergo much larger fluctuations (over a timescale of $\sim O(1000)$ rounds) than $|n_d|$; at this time I have no good explanation for this behaviour.

4.3 Individual agent behaviour in the VAEMG

4.3.1 Introduction

This section will focus on the behaviour of a single, arbitrary agent within the framework of the VAEMG model. I will examine three main properties of the model: 1. the performance of an agent with a variable level of activity; 2. an agent's *optimal* performance for a given set of model conditions; and 3. how this activity dependent behaviour varies with certain model parameters such as the number of agents, N, the inflation rate, I, and the bankruptcy level, Z_c.

How an agent's activity level affects his performance is, presently, the main concern. While the VAEMG is capable of, and, in fact, has been specifically designed to demonstrate this effect, the model includes unnecessary complexity that will detract from the *activity dependence* that is the present focus.

For this section I will use a simplified version of the VAEMG model that will provide a more direct window into the activity dependent properties of an individual agent. This model is identical to the full VAEMG with the exception that the gene value, p_i, is now the same for all agents ($p_i = p$), is fixed at $p = 1/2$, and is non-evolutionary (ie. $R_p = 0$).

The specific simplifications were chosen for two reasons. First, they reduce the complexity of the theory that was developed in section 3.3. While it might seem intuitive that eliminating the evolutionarity of p will simplify the associated theory, of almost equal importance in this simplification is the choice of $p = 1/2$; this point will be discussed further in section 4.3.3. Second, they simplify the dynamics of the game. In the second step of a round, each agent's strategy preference has been made irrelevant since he now has a $p = 1/2$ probability of choosing either to buy or sell. This second step is now equivalent to N_a agents each flipping an unbiased coin to determine their move for the round. By removing the evolutionarity of p we are eliminating the system's' dependence on its history so that we can more easily focus on the its α dependence.

4.3.2 Experimental results

Distribution of average lifetimes, $L(\alpha)$

Here I will focus on the 1-D distribution of average lifetimes, $L(\alpha)$, over the phase space of the VAEMG model.

Figure 4.2 shows $L(\alpha)$ for several values of the penalty, I. We can see from the figure that the system displays the proper behaviour at $\alpha = 0$, namely that $L(0) = Z_c/I$. Also, the distributions are all monotonic; for small values of $|I|$, $\frac{dL(\alpha)}{d\alpha} < 0$, while for larger values of $|I|$, $\frac{dL(\alpha)}{d\alpha} > 0$. When inflation is small an agent will do better by playing less frequently, but, eventually, as inflation is increased the agent will be better off by playing more. This dependence of $L(\alpha)$ on I will be discussed much further in later sections.

An immediate question that arises is whether or not the VAEMG reaches an equilibrium state and produces distributions that are stable with T. The model has been simulated for $T = 10^6 \rightarrow 10^9$ with various $I \in [-1, 0)$ and, ignoring fluctuations, in all cases $L(\alpha)$ is found to be stable in time (agreeing

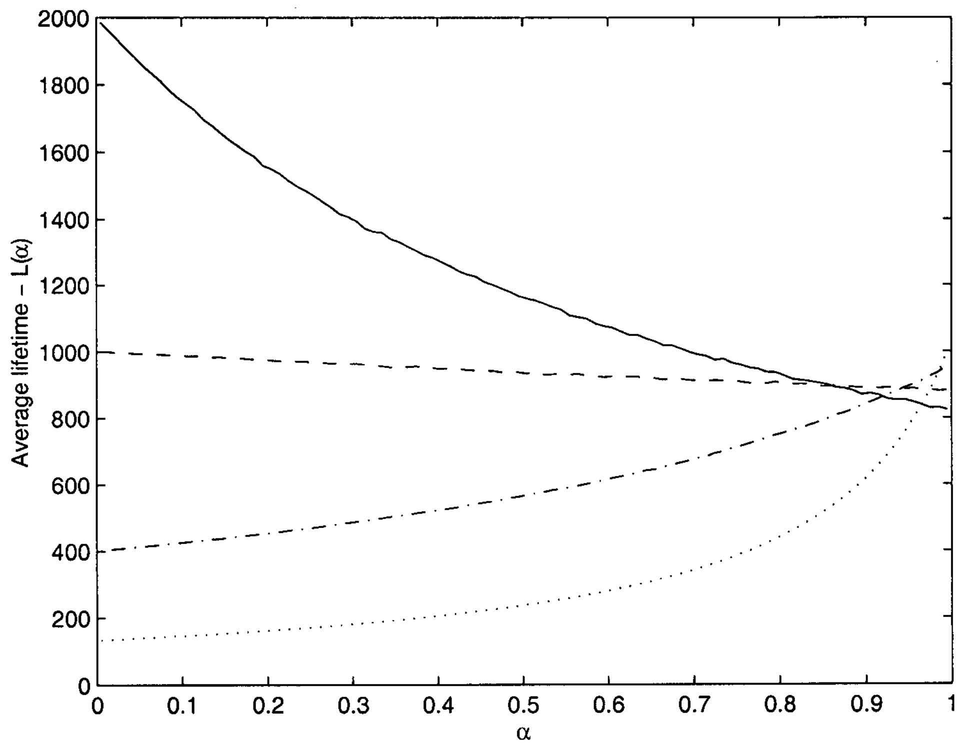

Figure 4.2: Average lifetime distributions for I = -0.05 (—), -0.10 (---), -0.25 (-·-·-), and -0.75 ($\cdots$) with $N = 101$, $Z_c = -100$, and $R_\alpha = 1$.

with the conclusions from above). With $I = 0$ this is not the case. Figure 4.3 shows $L(\alpha)$ at $I = 0$ normalized by $L_0(\alpha)$ (defined to be $L(\alpha)$ for $T = 10^6$) for several values of T. As T increases, the distribution becomes more heavily weighted around $\alpha = 0$. The reason for this result is that the state $\alpha = 0$ acts as a sink in the model. When an agent dies and is replaced by one with $\alpha = 0$, this new agent will live forever. As T increases, the probability of an agent attaining the $\alpha = 0$ state will increase; eventually (if T is large enough) all agents will become "trapped" in this sink and $L(\alpha) \to \delta(\alpha)$.

Figure 4.4 shows the variation in the distribution of average lifetimes with the number of agents, N. For all $\alpha \neq 0$, $L(\alpha)$ increases with increasing

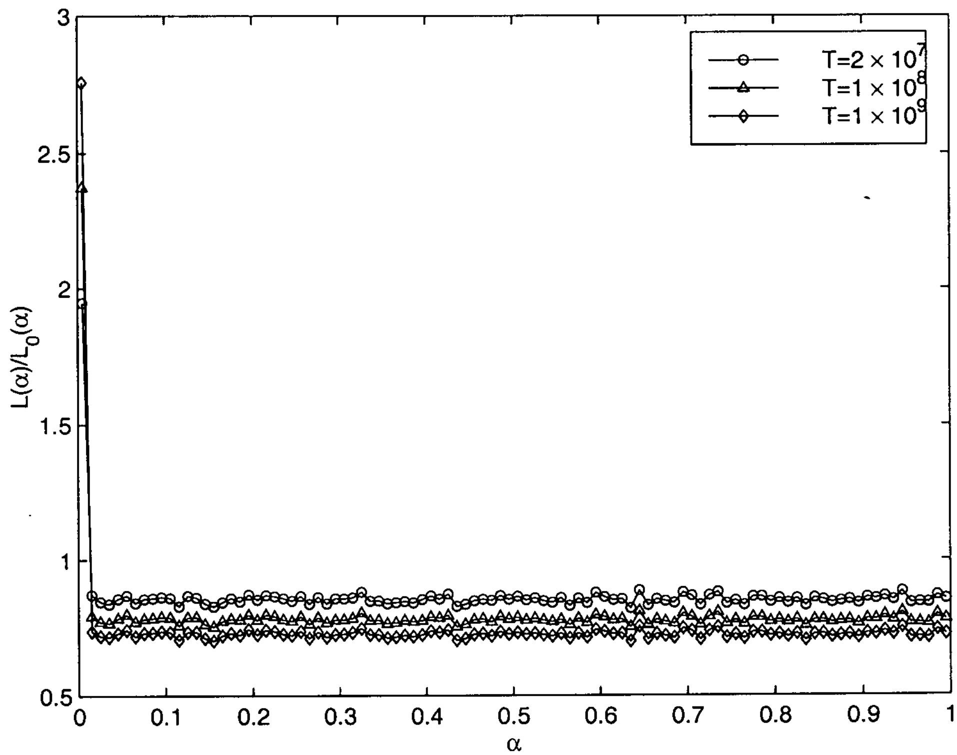

Figure 4.3: Normalized average lifetime distributions, $L(\alpha)/L_0(\alpha)$, for $T = 2 \times 10^7, 1 \times 10^8$ and 1×10^9 when the penalty, $I = 0$. Other parameters are $N = 101$, $Z_c = -100$, and $R_\alpha = 1$.

N, or, in other words, as the number of agents within the group increases, so too does the performance of each agent. To explain this result, time series of the number of active agents, N_a, have been plotted for several values of N in Figure 4.5. The second step of a round within the VAEMG is identical to the EMG with N_a agents. We know from chapter 2 that an agent's performance in the EMG increases with the number of agents. Since N_a increases with N (from figure 4.5), it is then clear why the performance also increases with N in the VAEMG.

In chapter 2 we saw that $L(p) \sim |Z_c|$. Since the the average point loss

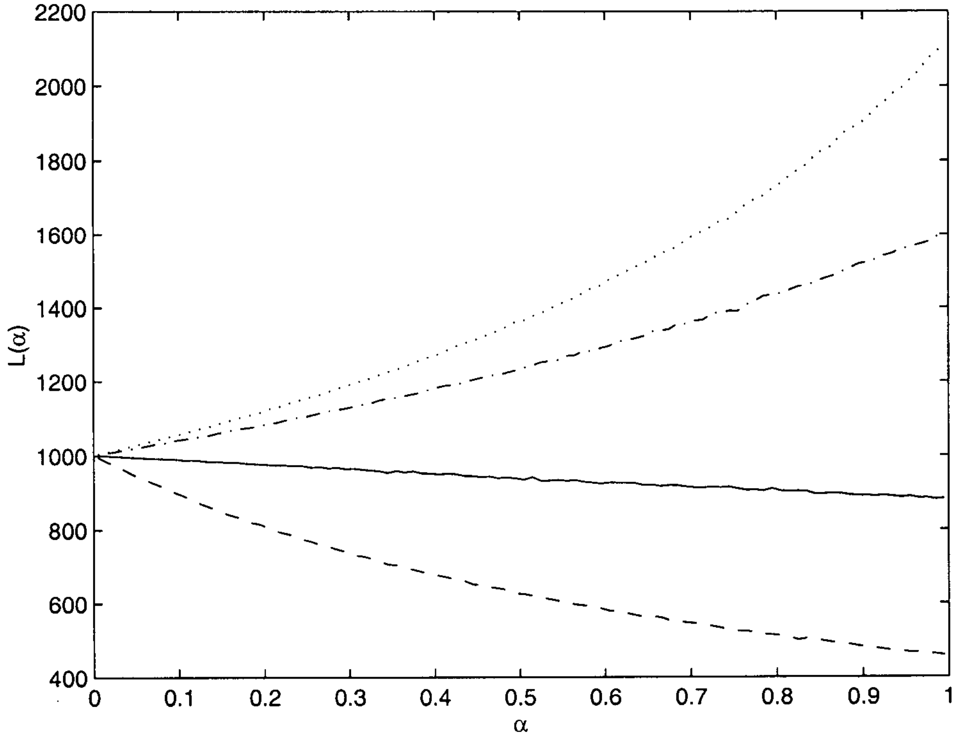

Figure 4.4: Average lifetime distributions for N=31 (---), 101 (—), 301 (-·-·-), and 501 ($\cdots$) when the penalty, I is -0.1. Other parameters are $Z_c = -100$ and $R_\alpha = 1$.

per round should be independent of the cutoff score, Z_c, it is reasonable to expect the same behaviour in the VAEMG. Figure 4.6 shows the dependence of average lifetime distribution on this "bankruptcy level", Z_c. While the scaling relation, $L(\alpha) \sim |Z_c|$, holds for large $|Z_c|$, it breaks down for $|Z_c| < 100$. Specifically, for $\alpha \neq 0$, an agent's average lifetime is larger than we expect in this small $|Z_c|$ region. The reason for this behaviour is due to the agent's inability to reach an equilibrium state. On average, we expect the ratio of number of active/inactive rounds to be approximately $\alpha/(1-\alpha)$ for each agent. When $|Z_c|$ is large, an agent's lifetime is large and this ratio is allowed to

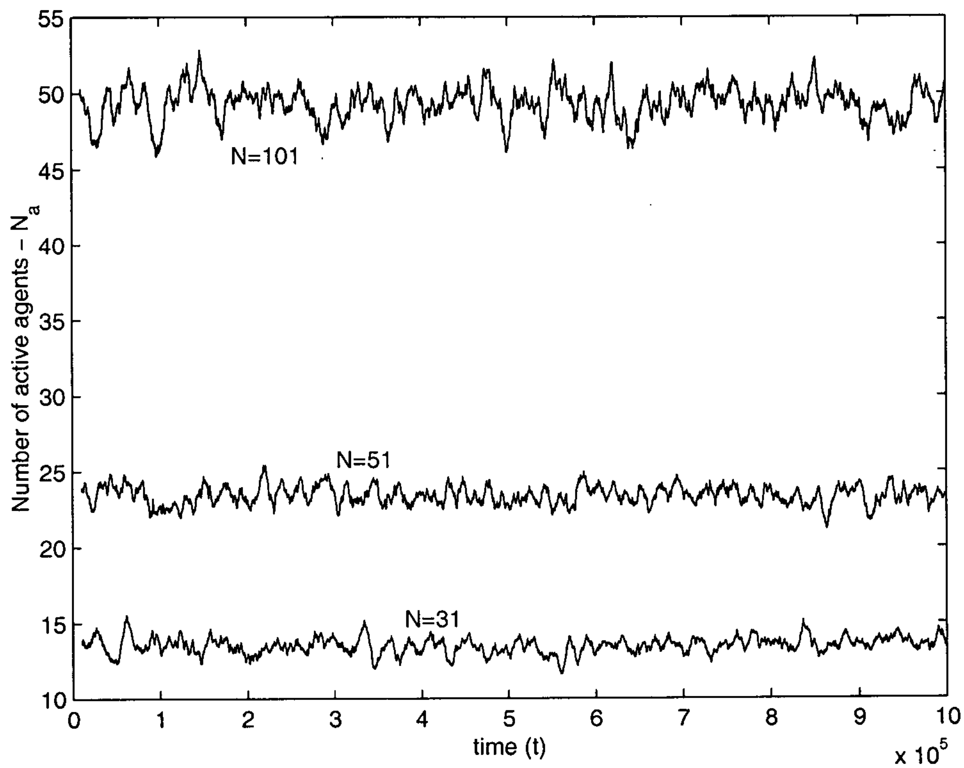

Figure 4.5: Time evolution of N_a for $N = 31, 51$, and 101 for $I = -0.1$. The time series have been averaged over a window of 10000 rounds.

approach its equilibrium value, but when $|Z_c|$ is small this will not usually happen. In this case, the rounds when an agent is inactive (and therefore loses only a small amount) become statistically more significant than they otherwise would be. Since the equilibrium behaviour of the VAEMG is the present focus I have used $Z_c = -100$ for all the simulations in this chapter.

Finally, simulations of the VAEMG have also been performed for various $R_\alpha \in (0, 1]$. For all values of this redistribution radius no variation was found in $L(\alpha)$.

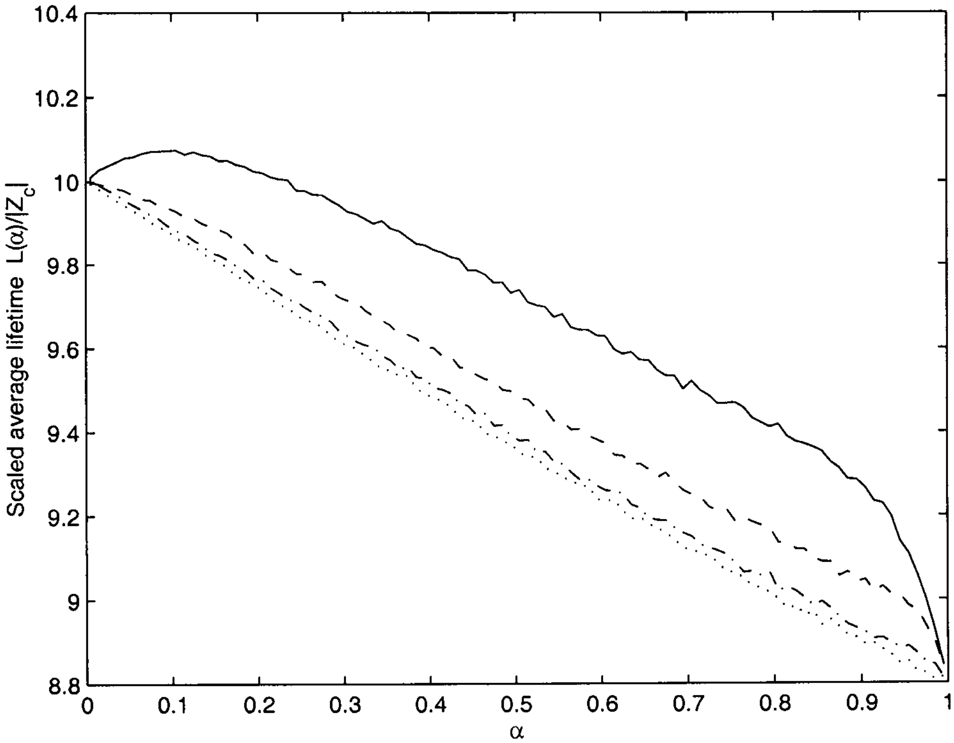

Figure 4.6: Average lifetime distributions for Z_c=-10 (—), -25 (---), -100 (-·-·-), and -200 (· · ·) with $I = -0.1$ and $N = 101$.

Optimal performance characteristics

In chapter 2, it was shown that an agent's optimal strategy is to either always or never follow the trend strategy, $\mathcal{S}^{\dagger}$; if the agent follows this rule then his average point loss per round will be minimized and he will survive the longest. The reason that the inactivity mechanism of the VAEMG has been modeled after the strategy selection mechanism of the EMG is so that we can determine an agent's optimal behaviour as a function of his activity level. To facilitate this study I will define α^* to be the activity level that results in an agent's optimal performance - ie. his maximal lifetime. For example, from figure 4.2,

$\alpha^*_{I=-0.05} = 0$ and so being inactive on every round will maximize the agent's lifetime.

Motivated by the results of figure 4.2, I have run simulations of the VAEMG with $N = 101$ for a range of I. From each distribution, $L(\alpha)$, the optimal activity level, α^*, has been determined and the results have been plotted vs. I in figure 4.7 below. Although not shown, $\alpha^* = 0\ \forall |I| < 0.1$ and

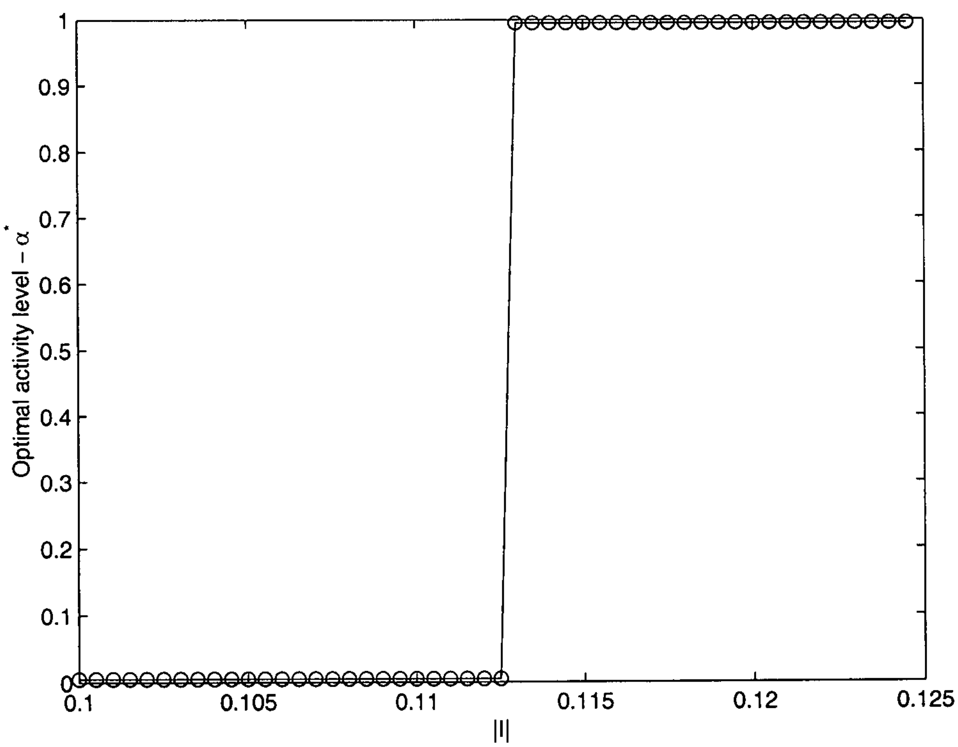

Figure 4.7: Dependence of the optimal activity level, α^*, on the inflation rate, $|I|$ for $N = 101$.

$\alpha^* = 1\ \forall |I| > 0.125$. These results show that there is a sharp transition from a pure strategy of *optimal inactivity* to the other pure strategy of *optimal activity* at some critical value of $|I^*| \approx 0.1125$. For $|I| < |I^*|$ it is optimal for an agent to remain inactive for every round of the game. If the agent participates in a round then there is a possibility that he will choose the minority option and increase his score, but, on average, his expected payoff will be less than if he had chosen not to play in the first place. As the rate of inflation is increased beyond $|I^*|$ the average loss being active will become smaller than I and so always playing the game will become the agent's optimal strategy.

Taking α^* to be the order parameter of the system suggests, due to the discontinuity in α^*, that the active/inactive transition is first order. To explore this point further, figure 4.8 shows several distributions, $L(\alpha)$, for $|I| \approx |I^*|$. Ignoring fluctuations, $L(\alpha)$ appears to be very linear near $|I^*|$ which

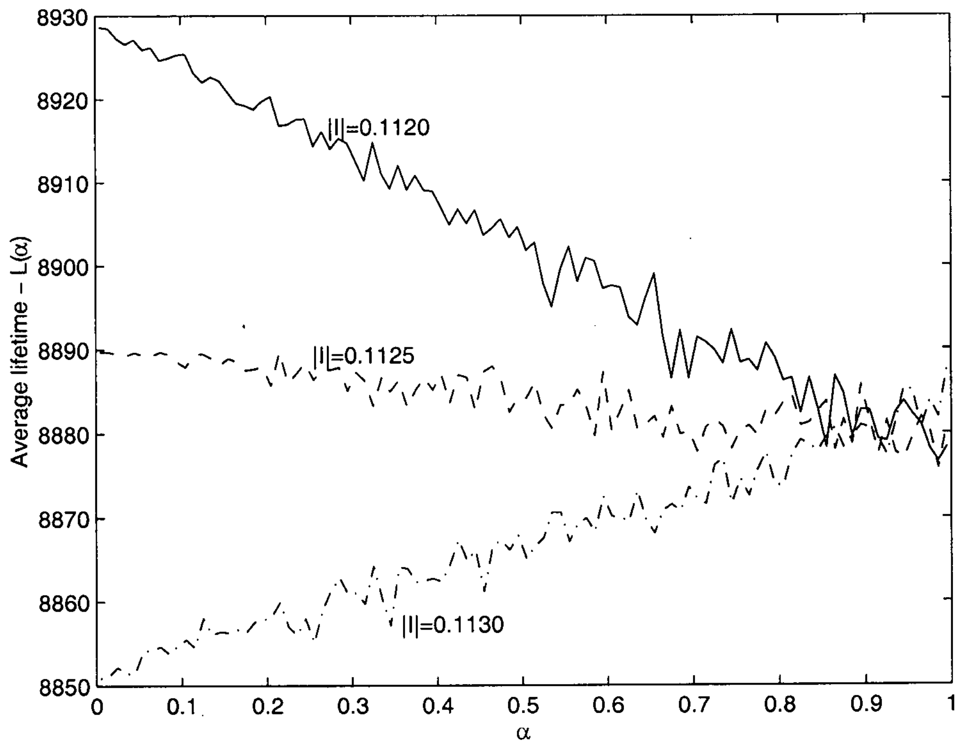

Figure 4.8: Average lifetime distributions, $L(\alpha)$, near the active/inactive phase transition, $|I^*|$, for $N = 101$.

confirms the first order nature of the transition. Furthermore, the linearity of

$L(\alpha)$ for $|I| \approx |I^*|$ suggests that *there is no preferred level of activity at the active/inactive transition*; on average, all values of α will result in an equal lifetime at $|I| = |I^*|$.

The distributions in figure 4.8 were run with $T = 10^9$. This was the minimal number of rounds required to reduce the size of the fluctuations in $L(\alpha)$ to the level where its maximum is distinguishable for $\Delta I = 0.0005$; this distinguishability is the reason why T must be so much larger for the VAEMG than for the EMG.

A note should be made about the nature of the optimality of α^*. For $|I| < |I^*|$, α^* is a true global optimum of this simplified ($p = 1/2$) version of the VAEMG. Regardless of the actions of the other $N-1$ agents in the group, when $|I| < |I^*|$ an agent's best strategy is to always remain inactive. When $|I| > |I^*|$, however, the strategy of $\alpha = 1$ (always play) is *only* optimal if the other agents play as they normally would. If a group of the background agents collude and decide to not play for a round then they can change the expected payoff of playing so that it is no longer the optimal strategy. The simplest example of this is the extreme case when $|I| > |I^*|$ but where all of the other $N-1$ agents "decide" not to play. Here the "optimal" strategy for the single agent under consideration is still to always play, but if he follows this strategy his expected payoff is -1 (a minimum) since he is guaranteed to always be in the majority group. This optimal strategy of activity is therefore not global since it is inherently dependent on the decisions of the other agents.

It will be useful to characterize the active/inactive phase transition in terms of the number of agents, N. Figure 4.9 shows plots of α^* vs. $|I^*|$ for various N. It is clear that the position of the transition, $|I^*|$, decreases with

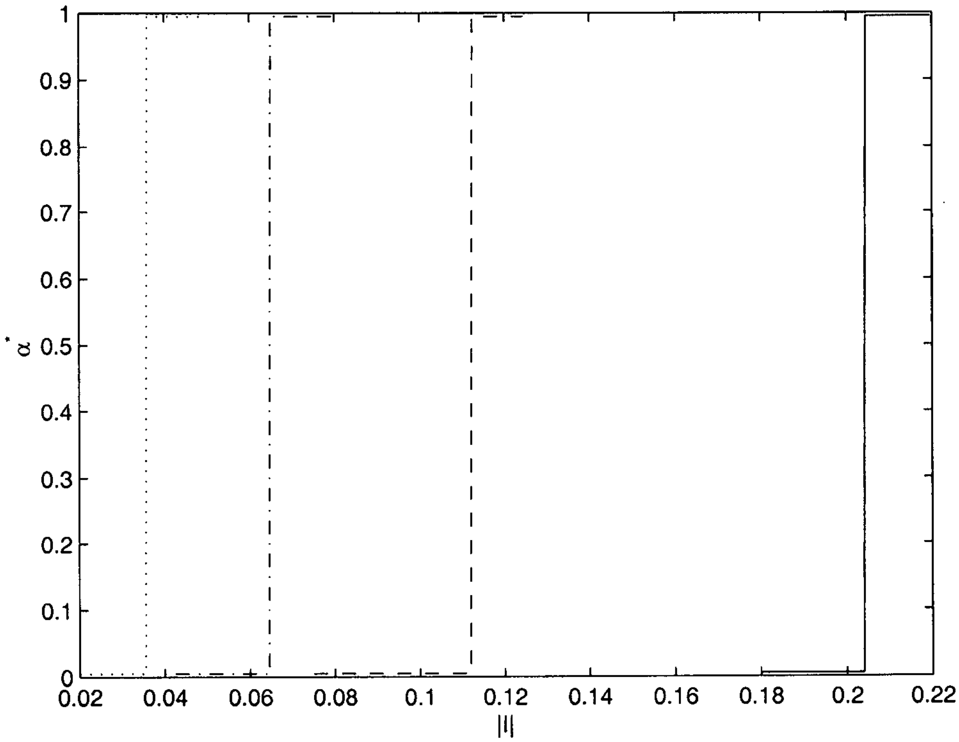

Figure 4.9: Active/inactive phase transitions for $N = 31$(—), 101(- - -), 301(-·-·), and 1001(· · ·).

increasing N. This result indicates that as the number of agents increases, the inflation rate that should "force" an agent into a state of activity decreases, or, in other words, an agent's optimal strategy is to take more risk as the size of the group increases.

4.3.3 Theoretical results

It was stated above that the structure of the simplified VAEMG used in this section reduces the complexity of the associated theory. Specifically, there are two main simplifications that result from the choice of $p_i = p = 1/2$. First, all functions of both α and p ($L(\alpha, p)$, $\langle S_c(\alpha, p)\rangle$, and $P_w(\alpha, p)$) now reduce to functions of α alone. As a result, the integration over p in Eq. (3.20) and (3.22) disappears and Eq. (3.21) reduces to

$$\mu_p = \frac{1}{2} \qquad\qquad \sigma_p^2 = 0 \tag{4.1}$$

To discuss the second simplification, I will restate Eq. (3.11) here.

$$p_w(N_a) = p_i \sum_{n=0}^{\frac{N_a-1}{2}} F_{N_a}(n) + (1-p_i) \sum_{n=\frac{N_a+1}{2}}^{N_a} F_{N_a}(n) - 2p_i(1-p_i) G^i_{N_a-1}\left(\frac{N_a-1}{2}\right) \tag{4.2}$$

Recall that $F_{N_a}(n)$ is the probability that n agents follow the trend strategy. Because of the choice of $p_i = 1/2$ and the fact that $\sum_{n=0}^{N_a} F_{N_a}(n) = 1$, Eq. (4.2) reduces to

$$p_w(N_a) = \frac{1}{2} - \frac{1}{2} G^i_{N_a-1}\left(\frac{N_a-1}{2}\right) \tag{4.3}$$

Since there is now no history dependence of the agents' strategy preferences, $G^i_{N_a-1}(n)$ also takes on a simplified form. The probability that n background agents choose the trend strategy is now simply the sum of N_a independent (resulting from the lack of a history dependence on p) decisions where n agents choose one of two possible states and $N_a - n$ agents choose the other. This probability can therefore be written as

$$\begin{aligned} G^i_{N_a-1} &= \frac{1}{2^n}\frac{1}{2^{N_a-n}}\binom{N_a}{n} \\ &= \frac{1}{2^{N_a}}\binom{N_a}{n} \end{aligned} \tag{4.4}$$

and so Eq. (4.3) becomes

$$p_w(N_a) = \frac{1}{2} - \frac{1}{2^{N_a+1}} \binom{N_a}{\frac{N_a-1}{2}} \tag{4.5}$$

The first term of Eq. (4.5) is simply an agent's probability of obtaining either of the two options when flipping a coin while the second is a finite size effect that reduces this probability given the games minority rule. The remainder of the theory is the same as was described in chapter 3.

Figure 4.10 shows a comparison of the experimental results of figure 4.2 with the theory of the simplified VAEMG. Agreement between the theory

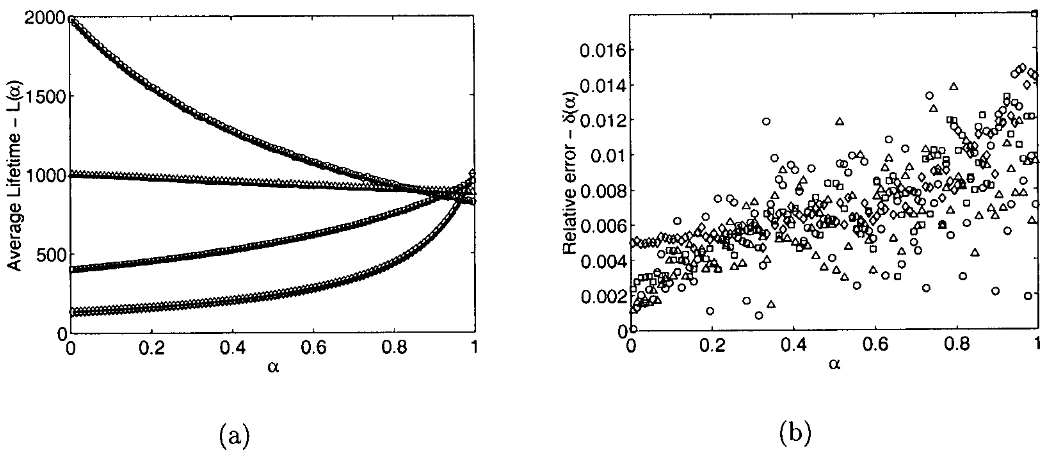

Figure 4.10: (a) Average lifetime distribution, $L(\alpha)$, for I = -0.05 (○), -0.1 (△), -0.25 (□), and -0.75 (◇) with N = 101. The solid lines are the respective theoretical distributions. (b) The relative error, $\delta(\alpha) = (L_{simulation}(\alpha) - L_{theory}(\alpha))/L_{simulation}(\alpha)$, of the theoretical distributions in (a).

and simulations is very good. Defining the relative error, $\delta(\alpha)$, of the the theoretical distribution as

$$\delta(\alpha) = \frac{L_{simulation}(\alpha) - L_{theory}(\alpha)}{L_{simulation}(\alpha)} \tag{4.6}$$

it can be seen from figure 4.10 (b) that $\delta(\alpha) < 2\%$ and much less than this in the majority of cases. $L_{theory}(\alpha)$ have been generated over the model's entire phase space (defined by N, I, and Z_c) and the results are in similar agreement as those above. Figure 4.10.b does indicate that $\delta(\alpha)$ is, in general, an increasing function of α and so is largest at $\alpha = 1$. Since the transition of α^* is most dependent on the relative height of $L(\alpha)$ at its endpoints, it is worth examining how well the theory predicts the position of this transition.

Figure 4.11 shows a comparison between the "experimental" active/ inactive phase transitions from section 4.3.2 and the theory. Again, the theory

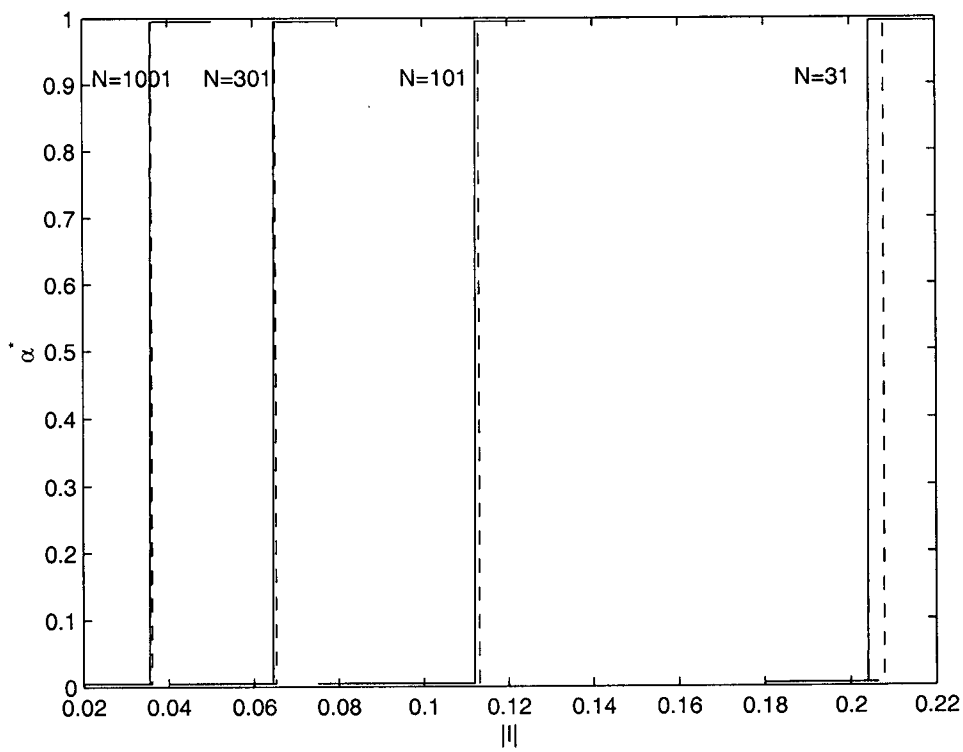

Figure 4.11: Comparison of "experimental" phase transition plots (—) (from figure 4.9) with theory (---) for $N = 31$, 101, 301, and 1001.

is in good agreement with the simulations and, moreover, becomes exact in the thermodynamic limit.

The discrepancy between theory and simulation is puzzling. Because of the simplified version of the model used in this section I am confident that Eq. (4.5) is exact. I am therefore left with the Gaussian assumption for $\Theta_N(N_a)$ as the only remaining source of error. To test this assumption $\Theta_N(N_a)$ has been measured from the simulated data and the results are shown in figure 4.12 below. The parabola-shaped curves in figure 4.12 demonstrate that, in

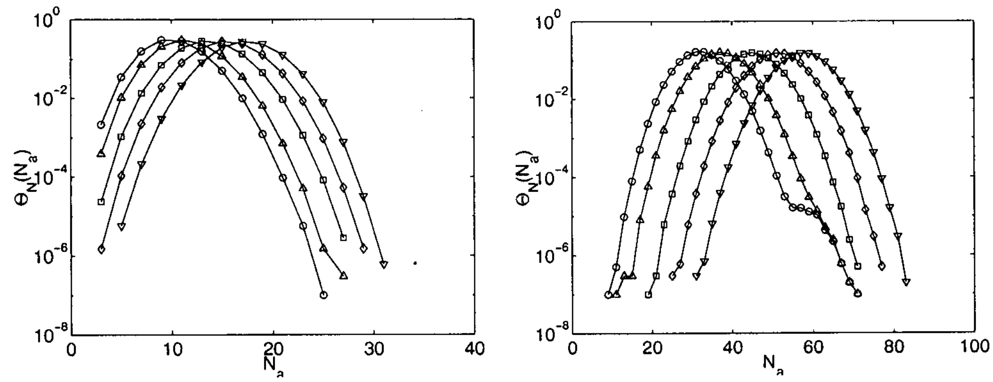

(a) N=31

(b) N=101

Figure 4.12: $\Theta_N(N_a)$ distributions for $N = 31$ (a) and $N = 101$ (b) for $I =$ -0.01($\circ$), -0.02 ($\triangle$), -0.05 ($\square$), -0.1 ($\diamond$), and -0.2 ($\bigtriangledown$).

most cases, the Gaussian assumption for $\Theta_N(N_a)$ appears justified. The only occasion when the assumption fails is for large N and small $|I|$ where the large N_a tail decays more slowly than a Gaussian. Since the disagreement between theory and simulation is most significant for small N (where the position of the transition, $|I^*|$, is largest) this failure in the assumption of a Gaussian $\Theta_N(N_a)$ is not able to account for the discrepancy.

This discrepancy, however, raises an interesting point about possible emergent behaviour within the VAEMG model. It was mentioned above that an agent's optimal level of risk increases as N increases; in figure 4.11 $|I|$ can therefore be thought of as a measure of this risk. What is interesting is that, for a given number of agents, N, an agent within the simulation acts as though he were less risk averse than a "theoretical" agent ($|I^*_{theory}| < |I^*_{simulation}|$). It should be noted that not only is α^* the *optimal* activity level for an agent, but it is also the agent's *most likely* activity level when the game is in equilibrium. Since the theory of the VAEMG has been developed using a purely mean-field approach, the more risk-preferring behaviour of an agent has possibly emerged through some more subtle mechanism generating correlations between agents other than the simple "mean-field" global history of the system.

4.4 Global behaviour of the VAEMG

In this section I will focus on the general VAEMG and concentrate on two results: 1. the behaviour of an agent within this generalized model; and 2. how the properties of the EMG are affected when the agents are given access to a third possible state. Due to the added complexity introduced by evolutionarity in p I will not attempt to characterize the model over its entire phase space. The purpose here will be to simply present a few selected results of the full VAEMG model that demonstrate the model's behaviour with respect to the above two categories.

Figure 4.13 shows the distribution of average lifetimes, $L(\alpha, p)$, generated by numerical simulation and the relative error of the corresponding theoretical distribution. The distribution demonstrates results that are con-

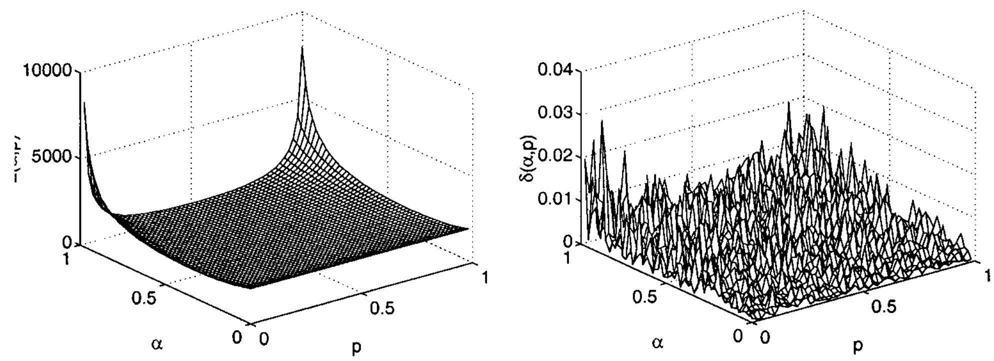

(a) (b)

Figure 4.13: (a) The distribution of average lifetimes, $L(\alpha, p)$, for $N = 101$, $I = -0.05$, $M = 3$, $Z_c = -100$, and $R_\alpha = R_p = 2$. (b) Relative error of the corresponding theoretical distribution.

sistent with those that have already been shown. At $\alpha = 1$, when an agent is always active, the distribution properly reduces to the bimodal distribution seen in figure 2.4. Also, although it is not easily seen in figure 4.13, at $p = 1/2$ the same convex, increasing nature of $L(\alpha, 1/2)$ (for $|I| < |I^*|$) that has been shown in section 4.3.2 is seen.

Theoretical distributions have been generated for a wide range of the model's phase space and agreement with the simulated data has been found to be very good in all cases (an example of which we have shown in figure 4.13 (b)). Because of this agreement, the discussion of the theoretical results in section 4.3.3 holds here as well and so no further mention will be made of the results of the theory.

A noticeable feature of figure 4.13 (a) is the apparent symmetry of $L(\alpha, p)$ with respect to p. To quantify this symmetry I will define a symmetry measure, η_α, for a fixed activity level, α, as

$$\eta_\alpha = \frac{\sum_{p<1/2} L(\alpha, p) - \sum_{p>1/2} L(\alpha, p)}{\sum_p L(\alpha, p)} \tag{4.7}$$

A plot of η_α for $I = -0.05$ is shown in figure 4.14. In chapter 2 we found that

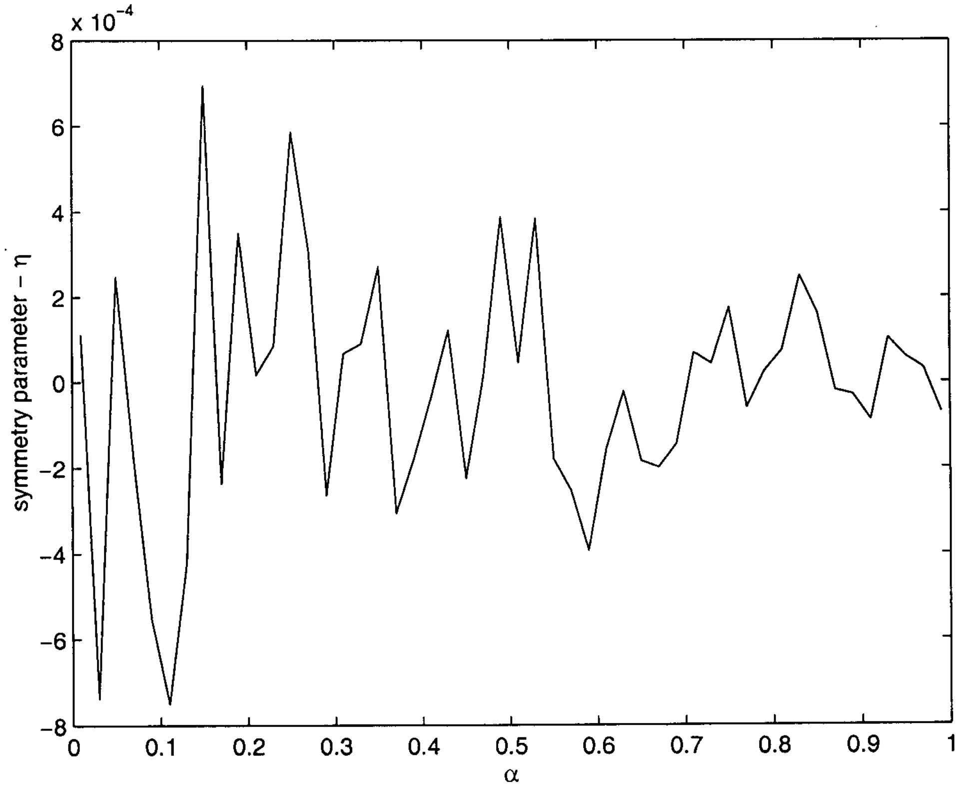

Figure 4.14: Measure of the p-symmetry in $L(\alpha, p)$ for the distribution shown in figure 4.13 (a).

an agent's optimal strategy is symmetric about $p = 1/2$ in the EMG. Ignoring fluctuations, this symmetry has been preserved in the VAEMG. With the exception of $\alpha = 0$ in which an agent's strategy preference, p, is irrelevant, for all α, $p = 0$ and $p = 1$ are optimal strategies. Furthermore, tossing a coin ($p = 1/2$) to determine whether to buy or sell is the worst possible strategy for any activity level, $\alpha > 0$.

Another noticeable feature of figure 4.13 is the optimal state for an agent in the full VAEMG. When evolutionarity in p is permitted, we can see that the system is able to access states that greatly outperform those when $p = 1/2$. To demonstrate the I-dependence of the state (α, p), figure 4.15 shows $L(\alpha, p)$ for several values of I. Together with figure 4.13 (a) these distributions

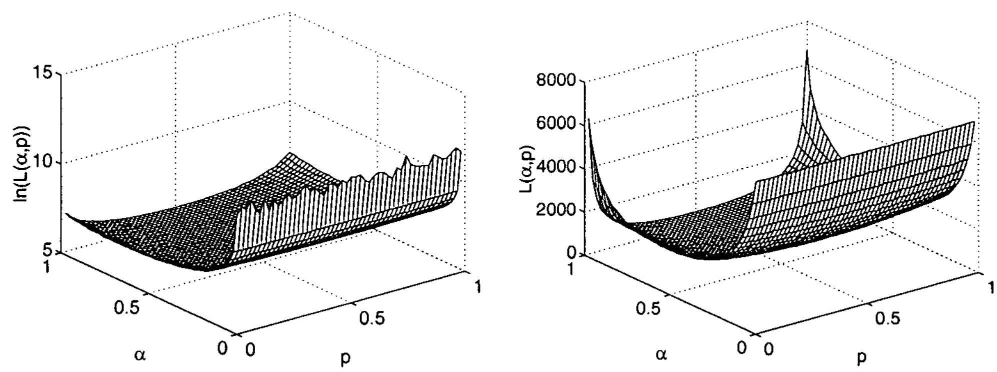

(a) (b)

Figure 4.15: The distribution of average lifetimes, $L(\alpha, p)$, for $I = 0.0$(a) and $I = -0.013$(b). Other parameters were $N = 101$, $Z_c = -100$, $M = 3$, and $R_\alpha = R_p = 2$.

demonstrate three interesting characteristics of the VAEMG's optimal state (α^*, p^*).

1. $|I^*| < |I^*_{p=1/2}|$. When an agent's gene value is able to evolve in response to the history of the game he should take more risk. For $N = 101$ the position of the phase transition has been determined to be $|I^*| = 0.0141 \pm 0.0001$.

2. For $|I| < |I^*|$ (see figure 4.13(a)) both $(\alpha^*, p^*) = (1{,}0)$ and $(1{,}1)$ are

equally optimal, while for $|I| > |I^*|$ (see figure 4.15(a)) all values of p are equally optimal (ignoring fluctuations).

3. For $|I| = |I^*|$ (see figure 4.15(b)), $L(\alpha, p)$ is a strictly convex function and so only the activities, $\alpha = 0$ and $\alpha = 1$ are optimal at this rate of inflation. This is in contrast to the $p = 1/2$ case where the distribution is α-independent at the transition and so all values of α are equally optimal.

There is another subtle, yet interesting property of the VAEMG that can be seen in the distributions of figure 4.15. Examining the optimal performance, $L(1, p^*)$, of an agent in the $\alpha = 1$ state shows that $L(1, p^*)$ is dependent on I. A more detailed plot of this relationship is shown in figure 4.16. Because the penalty, I, has no direct effect on the performance of an agent who is always active, to first approximation this result seems counterintuitive. But even though this agent is in the state $\alpha = 1$ there will always be other agents with $\alpha < 1$ and so, on average, $N_a < N$. Since an agent's performance increases with the number of agents playing the game (from figure 2.5 in chapter 2) and since N_a increases with $|I|$ it is therefore reasonable to expect that $L(1, p^*)$ to increase with $|I|$. Figure 4.16 shows this property for $0 \leq |I| \leq 0.1$, but for $|I| > 0.1$ the agent's optimal performance decreases with $|I|$. This large $|I|$ behaviour is contrary to the above reasoning and may be the result of some more subtle correlation between α and p.

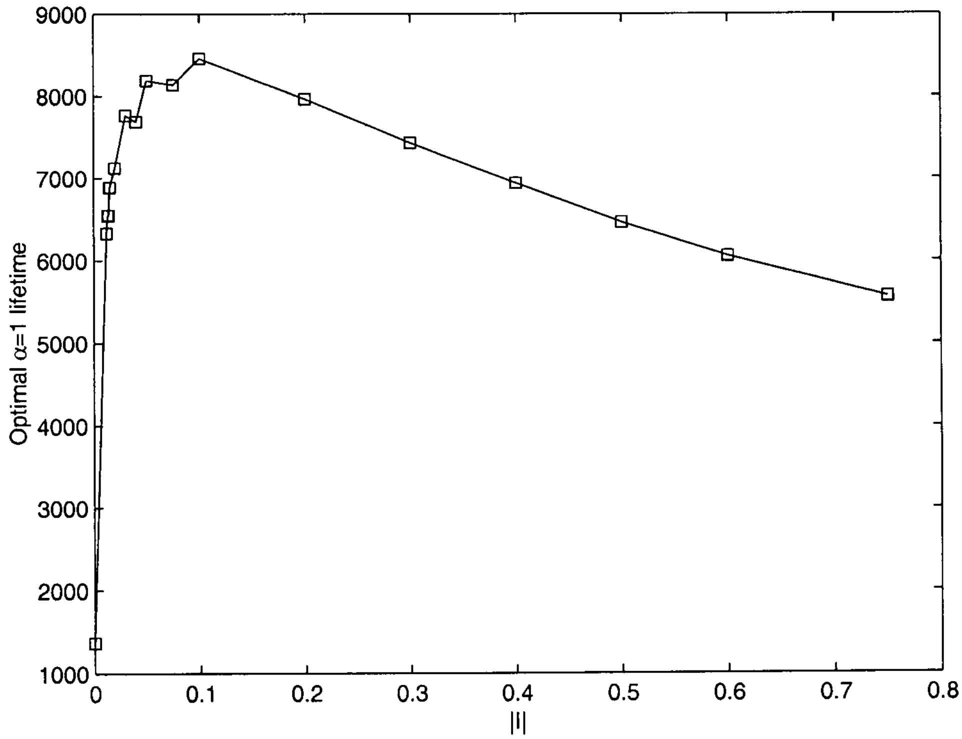

Figure 4.16: Optimal performance, $L(1, p^*)$, of an purely active agent vs. penalty, I. Other parameters are $N = 101$, $Z_c = -100$, $M = 3$, and $R_\alpha = R_p = 1/2$.

Chapter 5

Conclusions

5.1 Summary of results

In this paper I have developed an extension of the Evolutionary Minority Game (EMG), called the Variable Activity Evolutionary Minority Game (VAEMG), in which agents have the ability to sit-out on various rounds. The goal has been to characterize the model in terms of its parameters and determine how agent activity affects affects the behaviour of the game.

The original EMG model was designed to investigate how a group of agents compete for a limited resource when governed by supply-and-demand and a minority rule, both of which are believed to play important roles in the functioning of a financial market. Agent inactivity was the focus of this work for two main reasons. First, "forced" activity is one of the most unrealistic assumptions of the Minority Game. Inclusion of a variable activity level will serve to answer the question of how inactivity affects the specific structure of the EMG and will move us one step closer to the development of a more realistic market model. Second, from a purely academic standpoint, it is interesting to wonder how the inclusion of an asymmetric third state affects the dynamics of the self-organizing complex system formed by the EMG.

The structure of the inactivity mechanism in the VAEMG is intentionally simplistic. Work on the EMG has demonstrated how an agent's performance depends on his strategy preference, p; this preference, or *gene value*, is simply the probability that the agent will follow the trend in the game. In order to determine a similar relationship between the agent's performance and his activity level, the inactivity mechanism has been modeled in analogy with the gene value of the EMG. Finally, in order to retain the super-martingale nature of the EMG (ie. to prevent the trivial situation where an agent never participates, but lives forever), a penalty, which can be thought of as a form of inflation, has also been included in the model.

In order to focus on the activity dependence of an agent's performance, I first concentrated on a simplified version of the VAEMG. In this case each agent's gene value was fixed at $p = 1/2$ which was demonstrated to remove the system's history dependence on p to more easily focus on its activity-level dependence.

Many of the results of this simplified VAEMG can be understood directly in terms of the EMG. First, it has been found that the dynamics of the game are independent of an agent's memory size, M, and the redistribution radius, R_α, of his activity level. Also, with the exception of the special case, $I = 0$, where the $\alpha = 0$ state acts as a sink to the distribution and $L(\alpha) \to \delta(\alpha)$ as $T \to \infty$, an agent's behaviour has been found to be independent of the length of the game, T.

However, it has been found that $L(\alpha)$ does not scale with the bankruptcy level, Z_c, for all Z_c as with the EMG. For small $|Z_c|$ ($|Z_c| <\sim 100$) an agent tends to lose less per round than when $|Z_c|$ is large. This counter-intuitive result is caused by the non-integer nature of the penalty, I, and has been explained in terms of the system's inability to reach a fully equilibrated state when $|Z_c|$ is small.

The most significant result of the simplified VAEMG was that an agent is found to perform better (live longer) by playing more often as the inflation (measured by $|I|$) increases. Specifically, a first order transition from a phase of optimal inactivity to a phase of optimal activity has been found as $|I|$ is increased above some critical value, $|I^*|$. Finally, the position of this transition has been found to be strongly dependent on the number of agents within the group; specifically, $|I^*| \to 0$ in the thermodynamic limit.

Next, I studied the general formulation of the VAEMG in which both an agent's activity level and gene value were allowed to evolve in response to the history of the game; the purpose being to examine the global properties of the model and how the dynamics of the EMG are affected when the agents are allowed access to a third state.

Several interesting results have been observed with this model. First, it has been found that the VAEMG retains the symmetry in the gene value, p, as well as the characteristic bimodal dependence of an agent's performance on p. This result tells us that, for any given activity level, α, $p = 0$ and $p = 1$ are always an agent's equally optimal strategies. Furthermore, the optimal states

$(\alpha^*, p^*) = (1,0)$ and $(1,1)$ tend to dominate over all other $\alpha \neq 1$ states of the system even for values of the penalty much lower than the transition, $|I^*_{p=1/2}|$, of the simplified VAEMG. This result, that $|I^*_{general}| < |I^*_{p=1/2}|$, means that an agent's optimal strategy is to be much less risk averse when he is allowed to modify his strategy preference based on the history of the game. Finally, it has been found that the optimal performance of an agent who is always active, $L(1,p^*)$, is dependent on the value of I. This somewhat counterintuitive result has been explained in the region of small $|I|$ where $L(1,p^*)$ increases with $|I|$ using the empirical facts that the number of active agents, N_a, increases with $|I|$ and an agent's performance increases with N_a. For larger $|I|$, however, it has been found that $L(1,p^*)$ decreases with $|I|$ and it is believed that this is a result of some subtle correlation between an agent's activity level, α, and his gene value, p.

Finally, I have also developed an analytical theory of the VAEMG model based upon a mean-field approach. For both versions of the model (simplified and full) and for a large region of the system's phase space, the theory has been found to approximate the simulations very well and appears to become exact in the thermodynamic limit. Furthermore, the discrepancy at finite N has been explained in terms of an agent within the simulation taking on more risk than what is predicted by theory. Since the theory has been based on mean-field assumptions, this discrepancy indicates a possible emergence of subtle inter-agent correlations within the model for finite N.

5.2 Future directions

From its inception, the Minority Game was designed with a very simple structure to provide a solid base upon which to build a tractable model of a financial market. This paper has extended the basic EMG, but it is only one small step and there is still *a lot* of work left to be done.

Within the context of this paper, future work on the Minority Game will focus on two main areas. First, the structure of the inactivity mechanism is admittedly very simplistic in the VAEMG and, most likely, *too* simplistic to be used in a real-world model. It will be interesting study the VAEMG using inactivity rules that more closely resemble those found in a real market. For instance, while an agent's activity level, α, in the present model is dependent on the history of his performance, this dependence is only very weak; a much more realistic mechanism would incorporate specific rules based on past experience in the agent's decision making process. It will be interesting to see whether a more realistic mechanism is necessary to generate some of the stylized facts of a real market and, if so, what effect this will have on the global properties of the model.

Finally, as was stated above, many simplifying assumptions of the EMG are unrealistic and need modification before there is hope that the model can be used to make real-world predictions. Along with agent inactivity, work also needs to be done in developing realistic mechanisms for such things as the payoff structure, market maker, and diversity of the agent population. By formulating a market model from the ground up in this way hopefully we will be left with a model that is not only tractable, but also more realistic in its predictions.

Bibliography

[1] B. Mandelbrot. The variation of certain speculative prices. *Journal of Business*, 36:392–417, 1963.

[2] E. F. Fama. The behavior of stock-market prices. *Journal of Business*, 38:34–105, 1965.

[3] H.J. Blok. *On the nature of the stock market: Simulations and experiments.* PhD thesis, University of British Columbia, 2000. cond-mat/0010211.

[4] Jean-Phillipe Bouchaud and Marc Potters. *Theory of Financial Risks: From Statistical Physics to Risk Management.* Cambridge University Press, 2000.

[5] Rosario N. Mantegna and H. Eugene Stanley. *An Introduction to Econophysics: Correlations and Complexity in Finance.* Cambridge University Press, 1999.

[6] John Y. Campbell, Andrew W. Lo, and A. Craig MacKinlay. *The Econometrics of Financial Markets.* Princeton University Press, 1996.

[7] S. Galluccio, G. Caldarelli, M. Marsili, and Y.-C. Zhang. Scaling in currency exchange. *Physica A*, 245:423–436, 1997.

[8] R.N. Mantegna and H.E. Stanley. Scaling behavior in the dynamics of an economic index. *Nature*, 376:46–49, 1995.

[9] R.N. Mantegna and H.E. Stanley. Stock market dynamics and turbulence: Parallel analysis of fluctuation phenomenon. *Physica A*, 239:255–266, 1997.

[10] P. Cizeau, Y. Liu, M. Meyer, C.-K. Peng, and H.E. Stanley. Volatility distribution in the s&p 500 stock index. *Physica A*, 245:441–445, 1997.

[11] Y. Liu, P. Cizeau, M. Meyer, C.-K. Peng, and H.E. Stanley. Correlations in economic time series. *Physica A*, 245:437–440, 1997.

[12] Y. Liu, P. Gopikrishnan, P. Cizeau, M. Meyer, C.-K. Peng, and H.E. Stanley. Statistical properties of the volatility of price fluctuations. *Physical Review E*, 60:1390–1400, 1999.

[13] E.F. Fama. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25:383–417, 1970.

[14] T. Lux. Long-term stochastic dependence in financial prices: Evidence from the German stock market. *Applied Economics Letters*, 3:701–706, 1996.

[15] P. Gopikrishnan, M. Meyer, L.A.N Amaral, and H.E. Stanley. Inverse cubic law for the distribution of stock price variations. *The European Physical Journal B*, 3:139–140, 1998.

[16] P. Gopikrishnan, V. Plerou, Amaral L.A.N., M. Meyer, and H.E. Stanley. Scaling of the distribution of fluctuations of financial market indices. *Physical Review E*, 60:5305–5316, 1999.

[17] V. Plerou, P. Gopikrishnan, L.A.N. Amaral, M. Meyer, and H.E. Stanley. Scaling of the distribution of price fluctuations of individual companies. *Physical Review E*, 60:6519–6529, 1999.

[18] V. Plerou, P. Gopikrishnan, L.A.N. Amaral, X. Gabaix, and H.E. Stanley. Economic fluctuations and anomalous diffusion. *Physical Review E*, 62:3023–3026, 2000.

[19] S. Ghashghaie, W. Breymann, J. Peinke, P. Talkner, and Y. Dodge. Turbulent cascades in foreign exchange markets. *Nature*, 381:767–770, 1996.

[20] H.E. Stanley, P. Gopikrishnan, V. Plerou, and L.A.N. Amaral. Quantifying fluctuations in economic systems by adapting methods of statistical physics. *Physica A*, 287:339–361, 2000.

[21] A. Pagan. The econometrics of financial markets. *Journal of Empirical Finance*, 3:15–102, 1996.

[22] J. Hausman, A. Lo, and A.C. MacKinlay. An ordered probit analysis of stock transaction prices. *Journal of Financial Economics*, 31:319–379, 1992.

[23] R.R. Officer. The distribution of stock returns. *Journal of the American Statistical Association*, 67:807–812, 1972.

[24] P.D. Praetz. The distribution of share price changes. *Journal of Business*, 45:49–55, 1972.

[25] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 18:637–659, 1973.

[26] L. Bachelier. Théorie de la spéculation. *Annales Scientifiques de l'Ecole Normale Supérieure*, 3(17):21, 1900. Ph.D. Thesis.

[27] P. K. Clark. A subordinated stochastic process model with finite variance for speculative prices. *Econometrica*, 41:135–155, 1973.

[28] R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50:987–1007, 1982.

[29] T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327, 1986.

[30] R. N. Mantegna and H. E. Stanley. Modeling of financial data: Comparison of the truncated Lévy flight and the ARCH(1) and GARCH(1,1) processes. *Physica A*, 254:77–84, 1998.

[31] R. Cont and J.-P. Bouchaud. Herd behaviour and aggregate fluctuations in financial markets. *Macroeconomic Dynamics*, 4:170–196, 2000.

[32] A.-H. Sato and H. Takayasu. Dynamic numerical models of stock market price: from microscopic determinism to macroscopic randomness. *Physica A*, 250:231–252, 1998.

[33] H. Takayasu, T. Miura, T. Hirabayashi, and K. Hamada. Statistical properties of deterministic threshold elements. *Physica A*, 184:127–134, 1992.

[34] J. Masoliver, M. Montero, and J.M. Porrà. A dynamical model describing stock market price distributions. *Physica A*, 283:559–567, 2000.

[35] D. Sornette and A. Johansen. Large financial crashes. *Physica A*, 245:411–422, 1997.

[36] D. Sornette, A. Johansen, and J.-P. Bouchaud. Stock market crashes, precursors and replicas. *Journal de Physique I*, 6:167–175, 1996.

[37] A. Johansen and D. Sornette. Modeling the stock market prior to large crashes. *The European Physical Journal B*, 9:167–174, 1999.

[38] M. Youssefmir, B. Huberman, and T. Hogg. Bubbles and rarket crashes. *Computational Economics*, 12:97–114, 1998.

[39] D. Chowdhury and D. Stauffer. A generalized spin model of financial markets. *The European Physical Jounal B*, 8:477–482, 1999.

[40] J.-P. Bouchaud and R. Cont. A langevin approach to stock market fluctuations and crashes. *The European Physical Journal B*, 6:543–550, 1998.

[41] T. Lux and M. Marchesi. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397:497–500, 1999.

[42] F. Castiglione. Diffusion and aggregation in an agent based model of stock market fluctuations. *International Journal of Modern Physics C*, 11:865–879, 2000.

[43] T. Kaizoji. Speculative bubbles and crashes in stock markets: an interacting-agent model of speculative activity. *Physica A*, 287:493–506, 2000.

[44] R. Donangelo, A. Hansen, K. Sneppen, and S.R. Souza. Modelling an imperfect market. *Physica A*, 283:469–478, 2000.

[45] D.A. Vangheli and G. Ardelean. The ising like statistical models for studying the dynamics of the financial stock markets. cond-mat/0010318, 2000.

[46] M. Shatner, L. Mushnik, M. Leshno, and S. Solomon. A continuous time asynchronous model of the stock market; beyond the LLS model. cond-mat/0005430, 2000.

[47] S. Solomon. Generalized Lotka-Volterra (GLV) models. cond-mat/9901250, 1999.

[48] K.N. Ilinski and A.S. Stepanenko. Electrodynamical model of quasi-efficient financial market. cond-mat/9806138, 1998.

[49] D. Stauffer and D. Sornette. Self-organized percolation model for stock market fluctuations. *Physica A*, 271:496–506, 1999.

[50] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: An explanation of the $1/f$ noise. *Physical Review Letters*, 59:381–384, 1987.

[51] M. Paczuski, S. Maslov, and P. Bak. Avalanche dynamics in evolution, growth, and depinning models. *Physical Review A*, 53:414–443, 1996.

[52] S.E. Jorgensen, H. Mejer, and S.N. Nielsen. Ecosystem as self-organizing critical systems. *Ecological Modelling*, 111:261–268, 1998.

[53] H. Dawid. *Adaptive learning by genetic algorithms.* Springer Verlag, 1999.

[54] W.B. Arthur. Inductive reasoning and bounded rationality. *American Economic Review*, 84:406–411, 1994.

[55] D. Challet and Y.-C. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246:407–418, 1997.

[56] A. Cavagna, J.P. Garrahan, I. Giardina, and D. Sherrington. Thermal model for adaptive competition in a market. *Physical Review Letters*, 83:4429–4432, 1999.

[57] D. Challet, A. Chessa, M. Marsili, and Y.-C. Zhang. From minority games to real markets. cond-mat/0011042, 2001.

[58] D. Challet and Y.-C. Zhang. On the minority game: Analytical and numerical studies. *Physica A*, 256:514–532, 1998.

[59] D. Challet, M. Marsili, and Y.-C. Zhang. Modeling market mechanism with minority game. *Physica A*, 276:284–315, 2000.

[60] D. Challet and M. Marsili. Relevance of memory in minority games. *Physical Review E*, 62:1862–1868, 2000.

[61] M. Hart, P. Jefferies, N.F. Johnson, and P.M. Hui. Crowd-anticrowd theory of the minority game. cond-mat/0005152, 2000.

[62] P. Jefferies, M. Hart, P.M. Hui, and N.F. Johnson. From market games to real-world markets. cond-mat/0008387, 2000.

[63] P. Jefferies, M. Hart, N.F. Johnson, and P.M. Hui. Mixed population minority game with generalized strategies. *Journal of Physics A*, 33:L409–L414, 2000.

[64] N.F. Johnson, P.M. Hui, D.F. Zheng, and C.W. Tai. Minority game with arbitrary cutoffs. *Physica A*, 269:493–502, 1999.

[65] N.F. Johnson, P.M. Hui, D.F. Zheng, and M. Hart. Enhanced winnings in a mixed-ability population playing a minority game. *Journal of Physics A*, 32:L427–L431, 1999.

[66] Y. Li, A. VanDeemen, and R. Savit. The minority game with variable payoffs. *Physica A*, 284:461–477, 2000.

[67] Y. Li, R. Riolo, and R. Savit. Evolution in minority games. (I). Games with a fixed strategy space. *Physica A*, 276:234–264, 2000.

[68] Y. Li, R. Riolo, and R. Savit. Evolution in minority games. (II). Games with variable strategy spaces. *Physica A*, 276:265–283, 2000.

[69] M. Marsili and D. Challet. Trading behavior and excess volatility in toy markets. cond-mat/0004376, 2000.

[70] P.M. Hui, T.S. Lo, and N.F. Johnson. Segregation in a competing and evolving population. *Physica A*, 288:451–458, 2000.

[71] T.S. Lo, P.M. Hui, and N.F. Johnson. Theory of the evolutionary minority game. *Physical Review E*, 62:4393–4396, 2000.

[72] T.S. Lo, Lim S.W., Hui P.M., and N.F. Johnson. Evolutionary minority game with heterogeneous strategy distribution. *Physica A*, 287:313–320, 2000.

[73] Personal communications with Alan Kraus. Faculty of Commerce and Business Administration, University of British Columbia.

[74] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C: The art of scientific computing.* Cambridge University Press, second edition edition, 1992.