Evolution-Management in a Complex Adaptive System

David M.D. Smith
Lincoln College
University of Oxford

A thesis submitted for the degree of

PRS to DPhil

Summer 2003

Abstract

We consider the intriguing question of whether, by adjusting the heterogeneity of species in a complex adaptive system, one can manage and even control the future evolution in order to avoid specific undesired scenarios. We consider a simple yet highly non-trivial abstraction of such natural or man-made ecosystems related to the El Farol bar problem[2] and show that in principle such 'broad-brush' control is indeed possible. The intervention itself requires only minimal knowledge of the system's composition, and only involves a small number of microscopic changes, yet it serves to steer the subsequent global evolution away from undesirable regimes. This potential to provide global control by applying relatively minor perturbations at the level of the population's composition, seems quite remarkable given the large number of degrees of freedom in the system, and the inherent stochasticity both at the micro and macro level.

Contents

List of Figures

Chapter 1

Introduction

Complex Systems lie at the heart of many aspects of biology, engineering, computation and sociology. By their nature, such systems are very complicated to understand: they have many degrees of freedom corresponding to the large number of interacting objects or species within the system's population, the feedback and interactions lead to delicate temporal correlations, and there are stochastic influences both at the microscopic and macroscopic levels. Understanding emergent global behaviour is of great importance to the designers of such systems and is the focus of significant interest within the field of 'Collectives'[1]. For example, in a population consisting of many discreet control units [1], each of the units may itself malfunction whilst the population as a whole has to withstand global uncertainty due to interaction with the environment. Complete understanding and hence total control of such systems is impossible - instead one is aiming for some form of soft control whereby the probability that certain highly undesirable scenarios arise is minimised or even avoided. Even if in principle one knows what such scenarios are, and hence how to re-engineer the components such that they may be avoided, the necessary accessibility may not exist - for example, the individual units in a constellation of deployables may be too far away (for example, planetary rovers), or the system may have pre-built components which cannot be individually altered. In a more biological context, the approach to health-care whereby each cell of the body is re-designed is clearly impossible. Instead, more broad-brush approaches are required in which neither complete access to, nor information about, the individual components is assumed. We refer to [1] for a detailed discussion and perspective on the entire Collectives field.

[1] An example of such a population is that of trailing edge actuation devices used to suppress flutter in an aircraft wing[17]. Synchronising the actions of all of these edge effectors might be unfeasible for a human pilot.

In this report we address the question of whether, by adjusting the heterogeneity of species in a complex adaptive system, one can manage and even control the future evolution in order to avoid specific undesired scenarios. We consider a simple yet highly non-trivial binary abstraction of such natural or man-made ecosystems, called a B-A-R (Binary-Agent-Resource) system, and show that in principle such 'broad-brush' control is indeed possible. The intervention itself requires only minimal knowledge of the system's composition, and only involves a small number of microscopic changes, yet it serves to steer the subsequent global evolution away from undesirable regimes.

This potential to provide global control by applying relatively minor perturbations at the level of the population's composition, is far from obvious given the large number of degrees of freedom in the system, and the inherent stochasticity both at the macro and micro level (e.g. via the coin-tosses used to resolve tie-breaks in strategy scores). After all, the Central Limit Theorem tells us that the sum of a large number of stochastic processes of finite-variance will tend fairly rapidly to a Gaussian distribution. Hence, one might have thought that even with reasonably complete knowledge of the present and past states of the system, the evolution would be essentially diffusive and hence difficult to control without imposing substantial global constraints.

The report has several parts. We start by establishing a common framework for describing the spectrum of future paths of the complex adaptive system (Chapter 2). This framework is general to any complex system which can be mapped onto a general B-A-R (Binary Agent Resource) model in which the system's future evolution is governed by past history over an arbitrary but finite time window T (the $Time - Horizon$). This formalism then provides the tools whereby we can monitor the future evolution both with and without the perturbations to the population's composition. In Chapter 3, we discuss the B-A-R model in more detail, further information is provided in [16]. We emphasize that such B-A-R systems are *not* limited to the well-known El Farol Bar Problem [2] and Minority Games [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15] - instead these two examples are specific limiting cases. Initial investigations of a finite time-horizon version of the Minority Game were first presented in [15]. In Chapter 4, we consider the system's evolution in the absence of any such perturbations, hence representing the system's natural evolution. In Chapter 5, we revisit this evolution in the presence of control, where this control is limited to relatively minor perturbations at the level of the heterogeneity of the population. In Chapter 6 we discuss concluding remarks and possible extensions.

Chapter 2

Formal Description of the System's Evolution

Here we provide a general formalism applicable to any Complex System which can be mapped onto a population of N species or 'agents' who are repeatedly taking actions in some form of global 'game'. At each timestep each agent makes a (binary) decision $a_{\mu(t)}$ in response to the global information $\mu(t)$ which may reflect the history of past global outcomes. This global information is of the form of a bitstring of length m. For a general game, there exists some winning outcome $w(t)$ based on the aggregate action of the agents. Each agent holds a subset of all possible strategies - by assigning this subset randomly to each agent, we can mimic the effect of large-scale heterogeneity in the population. In other words, we have a simple way of generating a potentially diverse ecology of species, some of which may be similar but others quite different. One can hence investigate a typically-diverse ecology whereby all possible species are represented, as opposed to special cases of ecologies which may themselves generate pathological behavior due to their lack of diversity.

The aggregate action of the population at each timestep t is represented by $D(t)$, which corresponds to the accumulated decisions of all the agents and hence the (analog) output variable of the system at that timestep. The goal of the game, and hence the winning decision, could be to favour the minority group (MG), the majority group or indeed any function of the macroscopic or microscopic variables of the system. The individual agents do not themselves need to be conscious of the precise nature of the game, or even the algorithm for deciding how the winning decision is determined. Instead, they just know the global outcome, and hence whether their own strategies predicted the winning action[1]. The agents then reward the strategies in their posses-

[1]The algorithm used by the 'Game-master' to generate the winning decision, could also have a stochastic factor.

sion if the strategy's predicted action would have been correct if that strategy was implemented. The global history is then updated according to the winning decision. It can be expressed in decimal form as follows:

$$\mu(t) = \sum_{i=1}^{m} 2^{i-1}[w(t-i)+1] \ . \tag{2.1}$$

The system's dynamics are defined by the rules of the game. We will consider here the class of games whereby each agent uses his highest-scoring strategy at each timestep, and agents only participate if they possess a strategy with a sufficiently high success rate. [N.B. Both of these assumptions can be relaxed, thereby modifying the actual game being played]. The following two scenarios might then arise during the system's evolution:

- An agent has two (or more) strategies which are tied in score and are above the confidence level, and the decisions from them differ.
- The number of agents choosing each of the two actions is equal, hence the winning decision is undecided.

We will consider these cases to be resolved with a fair 'coin toss', thereby injecting stochasticity or 'noise' into the system's dynamical evolution. In the first case, each agent will toss his own coin to break the tie, while in the second the Game-master tosses a single coin. To reflect the fact that evolving systems will typically be non-stationary, and hence the more distant past will presumably be perceived as less relevant to the agents, the strategies are rewarded as to whether they would have made correct predictions over the last T timesteps of the game's running. There is no limit on the size of T other than it is finite and constant. The time-horizon represents a trajectory of length T on the de Bruijn graph in $\mu(t)$ (history) space[15] as shown in Figure 2.1. The stochasticity in the game means that for a given time-horizon T and a given strategy allocation in the population, the output of the system is not always unique. We will denote the set of all possible outputs from the game at some number of timesteps beyond the time-horizon T, as the Future-Cast (FC). It is useful to work in a time-horizon space $\underline{\Gamma_t}$ of dimension 2^{m+T}. An element Γ_t corresponds to the last $m+T$ elements of the bitstring of global outcomes (or equivalently, the winning actions) produced by the game. This dimension is constant in time whereas for a non-time-horizon game it would grow linearly. For any given time-horizon state, Γ_t, there exists a unique score vector $\underline{S(t)}$ which is the set of scores $S_R(t)$ for all the strategies which an agent could possess. As such, for each particular time-horizon

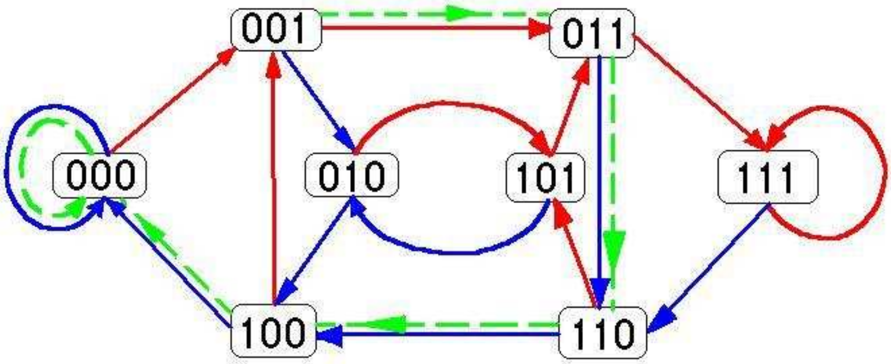

Figure 2.1: A path of time-horizon length $T = 5$ (dashed line) superimposed on the de Bruin graph for $m = 3$. The 8 global outcome states represent the 8 possible bit-strings for the global information, and correspond to the global outcomes for the past $m = 3$ timesteps.

state, there exists a unique probability distribution of the aggregate action, $D(t)$. This distribution of possible actions when a specified state is reached will necessarily be the same each time that state is revisited. Thus, it is possible to construct a transition matrix (c.f. Markov Chain) [15] $\underline{\underline{T}}$ of probabilities for the movements between these time-horizon states such that $\underline{P(\Gamma_t)}$ can be expressed as

$$\underline{P(\Gamma_t)} = \underline{\underline{T}} \; \underline{P(\Gamma_{t-1})} \tag{2.2}$$

where $\underline{P(\Gamma_t)}$ is a vector of dimension 2^{m+T} containing the probabilities of being in a given state Γ at time t

The transition matrix of probabilities is constant in time and necessarily sparse. For each state, there are only two possible winning decisions. The number of non-zero elements in the matrix is thus $\leq 2^{(m+T+1)}$. We can use the transition matrix in an eigenvector-eigenvalue problem to obtain the stationary state solution of $\underline{P(\Gamma)} = \underline{\underline{T}} \; \underline{P(\Gamma)}$. This also allows calculation of some time-averaged macroscopic quantities of the game [6][2].

[2]The steady state eigenvector solution is an exact expression equivalent to pre-multiplying the probability state vector $\underline{P(\Gamma_t)}$ by $\underline{\underline{T}}^{\infty}$. This effectively results in a probability state vector which is time-averaged over an infinite time-interval.

To generate the future-cast, we want to calculate the quantities in output space. To do this, we require;

- The probability distribution of $D(t)$ for a given time-horizon;
- The corresponding winning decisions, $w(t)$, for given $D(t)$;
- An algorithm generating output in terms of $D(t)$.

To implement the future-cast, we need to map from the transitions in the state space internal to the system to the macroscopic observables in the output space (often cumulative excess demand). We know that in the transition matrix, the probabilities represent the summation over a distribution of possible aggregate actions which is binomial in the case where the agents are limited to two possible decisions. Using the output generating algorithm, we can construct an 'adjacency' matrix $\underline{\underline{\Upsilon}}$ analogous to the transition matrix $\underline{\underline{T}}$, with the same dimensions. The elements of $\underline{\underline{\Upsilon}}$, contain probability distribution functions of change in output corresponding to the non-zero elements of the transition matrix together with the discrete convolution operator $\otimes$ whose form depends on that of the output generating algorithm. An explicit example of $\underline{\underline{\Upsilon}}$ is given in appendix A.

The adjacency matrix of functions and operators can then be applied to a vector, $\underline{\varsigma^{t=0}(S)}$, containing information about the current state of the game and of the same dimension as $\underline{\Gamma_t}$. $\underline{\varsigma^{t=0}(S)}$ not only describes the time-horizon state positionally through its elements but also the current value in the output quantity S within that element. At $t=0$, the state of the system is unique so there is only one non-zero element within $\underline{\varsigma^{t=0}(S)}$. This element corresponds to a probability distribution function of the current output value, its position within the vector corresponding to the current time-horizon state. The probability distribution function is necessarily of value unity at the current value or, for a future-cast expressed in terms of change in output from the current value, unity at the origin. The future-cast process for U timesteps beyond the present state can then be described by

$$\underline{\varsigma^{U}(S)} \;=\; \underline{\underline{\Upsilon}}^{U}\underline{\varsigma^{0}(S)}. \tag{2.3}$$

The actual future-cast, $\Pi(S,U)$, is then computed by superimposing the elements of the output/time-horizon state vector:

$$\Pi^{U}(S) = \sum_{i=1}^{2^{(m+T)}} \varsigma_i^{U}(S). \tag{2.4}$$

Thus the future-cast, $\Pi^U(S)$, is a probability distribution of the outputs possible at U timesteps in the future.

As a result of the state dependence of the Markov Chain, Π is non-Gaussian. As with the steady-state solution of the state space transition matrix, we would like to find a 'steady-state' equivalent for the output space[3] of the form

$$\Pi^1_{char}(S) = \left\langle \Pi^1(S) \right\rangle_\infty \tag{2.5}$$

where the one-timestep future-cast is time-averaged over an infinitely long period. Fortunately, we have the steady state solutions of $\underline{P(\Gamma)} = \underline{\underline{T}}\ \underline{P(\Gamma)}$ which are the (static) probabilities of being in a given time-horizon state at any time. By representing these probabilities as the appropriate functions, we can construct an 'initial' vector, $\underline{\kappa}$, similar in form to $\underline{\varsigma(S,0)}$ in (2.3) but equivalent to the eigenvector solution of the Markov Chain. We can then generate the solution of (2.5) for the *characteristic* future-cast, Π^1_{char}, for a given initial set of strategies. An element κ_i is again a probability distribution which is simply the point $(0\ ,\ P_i(\Gamma))$, the static probability of being in the time-horizon state denoted by the elements position, i. We can then get back to the future-cast

$$\Pi^1_{char}(S) \;=\; \sum_{i=1}^{2^{(m+T)}} \varsigma^1_i \quad \text{where} \quad \underline{\varsigma^1} \;=\; \underline{\underline{\Upsilon}}\ \underline{\kappa}. \tag{2.6}$$

We can also generate characteristic future-casts for any number of timesteps, U, by pre-multiplying $\underline{\kappa}$ by $\underline{\underline{\Upsilon}}^U$

$$\Pi^U_{char}(S) \;=\; \sum_{i=1}^{2^{(m+T)}} \varsigma^U_i \quad \text{where} \quad \underline{\varsigma^U} \;=\; \underline{\underline{\Upsilon}}^U\ \underline{\kappa}\ . \tag{2.7}$$

We note that Π^U_{char} is not equivalent to the convolution of Π^1_{char} with itself U times and as such is not necessarily Gaussian. The characteristic future-cast over U timesteps is simply the future-cast of length U from all the 2^{m+T} possible initial states where each contribution is given the appropriate weighting factor. This factor corresponds to the probability of being in that initial state. The characteristic future-cast can also be expressed as

$$\Pi^U_{char}(S) \;=\; \sum_{\Gamma=1}^{2^{(m+T)}} P(\Gamma)\ \Pi^U(S) \mid \Gamma \tag{2.8}$$

where $\Pi^U(S) \mid \Gamma$ is a normal future-cast from an initial time-horizon state Γ and $P(\Gamma)$ is the static probability of being in that state at a given time.

[3]Note that we can use this framework to generate time-averaged quantities of *any* of the macroscopic quantities of the system (e.g total number of agents playing) or volatility (Appendix B).

Chapter 3

Binary Agent Resource (B-A-R) system

The general binary framework of the B-A-R (Binary Agent Resource) system was discussed in Chapter 2. The global outcome of the 'game' is represented as a binary digit which favours either those choosing option +1 or option −1 (or equivalently 1 or 0, A or B etc.). The agents are randomly assigned S strategies at the beginning of the game. Each strategy comprises an action $a^s_{\mu(t)}$ in response to each of the 2^m possible histories μ, thereby generating a total of 2^{2^m} strategies in the Full Strategy Space [1]. At each turn of the game, the agents employ their most successful strategy, being the one with the most virtual points. The agents are thus adaptive if $S > 1$. We have already extended the B-A-R system by introducing the time-horizon T, which determines the number of past timesteps over which virtual points are collected for each strategy. We further extend the system by the introduction of a confidence level. The agents decide whether to participate or not depending on the success of their strategies. As such, the number of active agents $N(t)$ is less than or equal to N_{tot} at any given timestep. This results in a variable number of participants per timestep $V(t)$, and constitutes a 'Grand Canonical' game. The threshold, τ, denotes the confidence level: each agent will only participate if he has a strategy with at least r points where

$$r = T(2\tau - 1). \tag{3.1}$$

Agents without an active strategy become temporarily inactive.

[1] We note that many features of the game can be reproduced using a Reduced Strategy Space of 2^{m+1} strategies, containing strategies which are either anti-correlated or uncorrelated with each other[3]. The framework established in the present paper is general to both the full and reduced strategy spaces, hence the full strategy space will be adopted here.

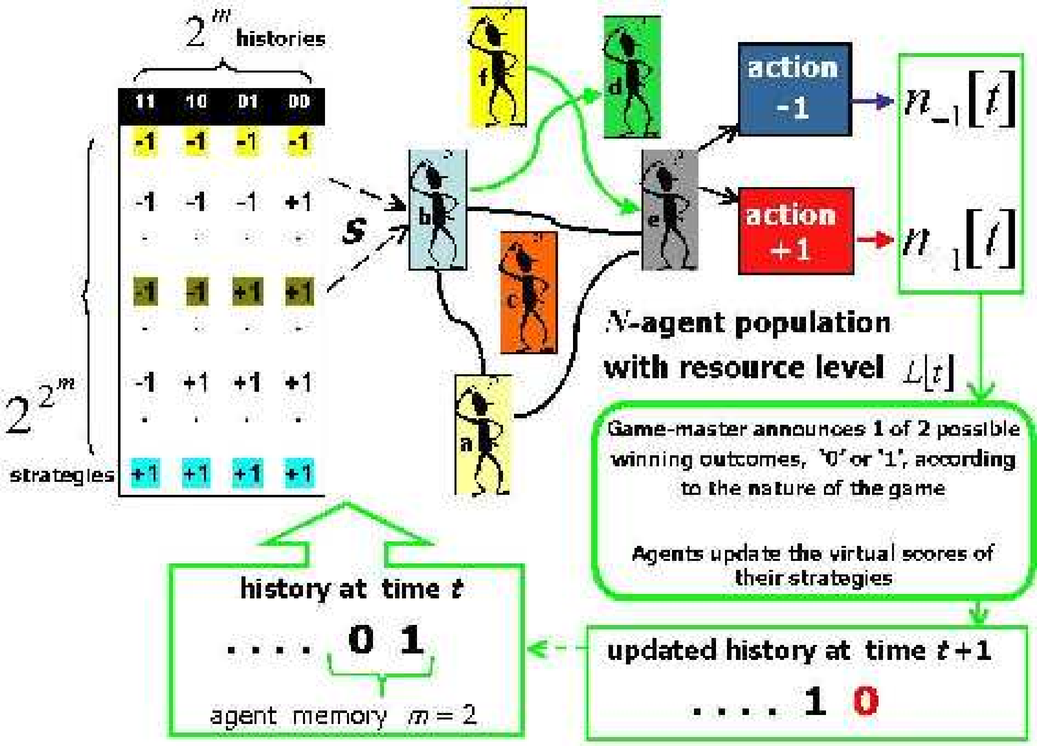

Figure 3.1: Schematic diagram of the Binary Agent Resource (B-A-R) system.

In keeping with typical biological, ecological, social or computational systems, the Game-master takes into account a finite global resource level when deciding the winning decision at each timestep. For simplicity, we will here consider the specific case whereby the resource level $L(t) = \phi V(t)$ with $0 \leq \phi \leq 1$[2]. We denote the number of agents choosing action +1 (or equivalently A) as $N_{+1}(t)$, and those that choose action -1 (or equivalently B) as $N_{-1}(t)$. If $L(t) - N_{+1}(t) > 0$ the winning action is +1 and vice-versa. We define the winning decision 1 or 0 as follows:

$$w(t) = \text{step}[L(t) - N_{+1}(t)] \tag{3.2}$$

where we define step$[x]$ to be

$$\text{step}[x] = \begin{cases} 1 & \text{if } x > 0, \\ 0 & \text{if } x < 0, \\ \text{fair coin toss} & \text{if } x = 0. \end{cases} \tag{3.3}$$

When $x = 0$, there is no definite winning option since $N_{+1}(t) = N_{-1}(t)$, hence the Game-master uses a random coin-toss to decide between the two possible outcomes.

[2]We note that ϕ itself could be actually be a stochastic function of the known system parameters.

We use a binary payoff rule for rewarding strategy scores, although more complicated versions can, of course, be used. However, we note that non-binary payoffs (e.g. a proportional payoff scheme) will decrease the probability of tied strategy scores, hence making the system more deterministic. Since we are interested in seeing the extent to which stochasticity can prevent control, we are instead interested in preserving the presence of such stochasticity. The reward function χ can be written

$$\chi[N_{+1}(t), L(t)] = \begin{cases} 1 & \text{for } w(t) = 1, \\ -1 & \text{for } w(t) = 0, \end{cases} \tag{3.4}$$

namely +1 for predicting the correct action and -1 for predicting the incorrect one. For a given strategy, R, the virtual points score is given by

$$S_R(t) = \sum_{i=t-T}^{t-1} a_R^{\mu(i)} \chi[N_{+1}(i), L(i)], \tag{3.5}$$

where $a_R^{\mu(t)}$ is the response of strategy, R, to the global information $\mu(t)$ summed over the rolling window of width T. The global output signal $D(t) = N_{+1}(t) - N_{-1}(t)$ is calculated at each iteration to generate an output time series.

Chapter 4

Natural Evolution: no system management

To realize all possible paths within a given game is necessarily computationally expensive. For a future-cast U timesteps beyond the current game state, there are necessarily 2^U winning decisions to be considered. Fortunately, not all winning decisions are realized by a specific game and the numerical generation of the future cast can be made reasonably efficient (see later).

Fortunately we can approach the future cast analytically *without* having to keep track of the agents' individual microscopic properties. Instead we group the agents together via the population tensor of rank S given by $\underline{\underline{\Omega}}$, which we will refer to as the Quenched Disorder Matrix (QDM) [8]. This matrix is assumed to be constant over the time-scales of interest, and more typically is fixed at the beginning of the game. The entry $\Omega_{R2,R2,\ldots}$ represents the number of agents holding the strategies $R1, R2, \ldots$ such that

$$\sum_{R,R',\ldots} \underline{\underline{\Omega}}_{R,R',\ldots} = N \tag{4.1}$$

For numerical analysis, it is useful to construct a symmetric version of this population tensor, $\underline{\underline{\Psi}}$. For the case $S = 2$, we will let $\underline{\underline{\Psi}} = \frac{1}{2}(\underline{\underline{\Omega}} + \underline{\underline{\Omega}}^T)$ [5].

The output variable $D(t)$ can be written in terms of the decided agents $D_D(t)$ who act in a pre-determined way since they have a unique predicted action from their strategies, and the undecided agents $D_U(t)$ who require an additional coin-toss in order to decide which action to take. Hence

$$D(t) = D_D(t) + D_U(t). \tag{4.2}$$

We focus on $S = 2$ strategies per agent although the approach can be generalized. The element $\Psi_{R,R'}$ represents the number of agents holding both strategy R and R'.

We can now write $D_D(t)$ as

$$D_D(t) = \sum_{R=1}^{Q} a_R^{\mu(t)} \mathcal{H}[S_R(t) - r] \sum_{R'=1}^{Q} (1 + \mathrm{sgn}[S_R(t) - S_{R'}(t)])\Psi_{R,R'} \tag{4.3}$$

where Q is the size of the strategy space, $\mathcal{H}$ is the Heaviside function and sgn$[x]$ is defined as

$$\mathrm{sgn}[x] = \begin{cases} 1 & \text{if } x > 0, \\ -1 & \text{if } x < 0, \\ 0 & \text{if } x = 0. \end{cases} \tag{4.4}$$

The volume $V(t)$ of active agents can be expressed as

$$V(t) = \sum_{R,R'} \mathcal{H}[S_R(t) - r]\big\{sgn[S_R(t) - S_{R'}(t)] + \frac{1}{2}\delta[S_R(t) - S_{R'}(t)]\big\}\Psi_{R,R'} \tag{4.5}$$

where δ is the Dirac delta. The number of undecided agents $N_U(t)$ is given by

$$N_U(t) = \sum_{R,R'} \mathcal{H}[S_R(t) - r]\delta(S_R(t) - S_{R'}(t))[1 - \delta(a_R^{\mu(t)} - a_{R'}^{\mu(t)})]\Psi_{R,R'}. \tag{4.6}$$

We note that for $S = 2$, each undecided agent's contribution to $D(t)$ is an integer, hence the demand of all the undecided agents $D_U(t)$ can be written simply as

$$D_U(t) \ \epsilon \ 2\mathrm{Bin}\bigg(N_U(t), \frac{1}{2}\bigg) - N_U(t) \tag{4.7}$$

where Bin(n, p) is a sample from a binomial distribution of n trials with probability of success p.

For any given time-horizon space-state Γ_t, the score vector $\underline{S(t)}$ (i.e., the set of scores $S_R(t)$ for all the strategies in the QDM) is unique. Whenever this state is reached, the quantity $D_D(t)$ will necessarily always be the same, as will the distribution of $D_U(t)$. We can now construct the transition matrix $\underline{\underline{T}}$ giving the probabilities for the movements between these time-horizon states. The element $\underline{\underline{T}}_{\Gamma_t|\Gamma_{t-1}}$ which corresponds to the transition from state Γ_{t-1} to Γ_t, is given for the (generalizable) $S = 2$ case by

$$\begin{aligned} \underline{\underline{T}}_{\Gamma_t|\Gamma_{t-1}} = \sum_{x=0}^{N_U} \Bigg\{ {}^{N_U}C_x(\frac{1}{2})^{N_U} \delta\Bigg[Sgn(D_D + 2x - N_U + V(1 - 2\phi)) \ + \\ (2\mu_t\%2 - 1)\Bigg] \ + \\ {}^{N_U}C_x(\frac{1}{2})^{(N_U+1)} \delta\Bigg[Sgn(D_D + 2x - N_U + V(1 - 2\phi)) \ + \ 0\Bigg] \Bigg\} \end{aligned} \tag{4.8}$$

where N_U, D_D implies $N_u(\Gamma_{t-1})$ and $D_D(\Gamma_{t-1})$, V implies $V(t-1)$, ϕ sets the resource level as described earlier and $\mu_t\%2$ is the required winning decision to get from state Γ_{t-1} to state Γ_t. We use the transition matrix in the eigenvector-eigenvalue problem to obtain the stationary state solution of $\underline{P(\Gamma)} \;=\underline{\underline{T}}\;\underline{P(\Gamma)}$. The probabilities in the transition matrix represent the summation over a distribution which is binomial in the $S=2$ case. These distributions are all calculated from the QDM which is fixed from the outset. To transfer to output-space, we require an output generating algorithm. Here we use the equation

$$S(t+1) = S(t) + D(t) \tag{4.9}$$

hence the output value $S(t)$ represents the cumulative value of $D(t)$, while the increment ΔS is simply $D(t)$. We define the operator $\otimes$ is as

$$(f \otimes g)\mid_i \;\; = \sum_{j=-\infty}^{\infty} f(i-j) \times g(j). \tag{4.10}$$

The formalism could be extended for general output algorithms using differently defined convolution operators.

An element in the adjacency for the $S=2$ case can then be expressed as

$$\begin{aligned} \underline{\underline{\Upsilon}}_{\Gamma_t|\Gamma_{t-1}} = \Bigg\{ \sum_{x=0}^{N_U} \Bigg(&(D_D + 2x - N_U), \\ {}^{N_U}C_x(\frac{1}{2})^{N_U} \delta \Big[&Sgn(D_D + 2x - N_U + V(1-2\phi)) + (2\mu_t\%2 - 1) \Big] + \\ {}^{N_U}C_x(\frac{1}{2})^{(N_U+1)} \delta \Big[&Sgn(D_D + 2x - N_U + V(1-2\phi)) + 0 \Big] \Bigg) \Bigg\} \otimes \end{aligned} \tag{4.11}$$

where N_U, D_D again implies $N_u(\Gamma_{t-1})$ and $D_D(\Gamma_{t-1})$, V implies $V(t-1)$, and $\mu_t\%2$ is the winning decision necessary to move between the required states. The future-cast and characteristic future-casts ($\Pi(S, U)$, Π_U) can then be computed for a given initial quenched disorder matrix (QDM).

We now consider an example to illustrate the implementation. In particular, we provide the explicit solution of a future-cast in the regime of small m and T, given the randomly chosen quenched disorder matrix

$$\underline{\underline{\Omega}} = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 2 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 2 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 2 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 2 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 2 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 1 \\ 2 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 3 & 1 & 1 & 0 & 1 & 0 & 1 \\ 0 & 2 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 3 & 2 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 2 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 2 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 \end{pmatrix}. \tag{4.12}$$

We consider the full strategy space and the the following game parameters:

Number of agents	N_{tot}	101
Memory size	m	2
Strategies per agent	s	2
Resource level	ϕ	0.5
Time horizon	T	2
Threshold	τ	0.51

The dimension of the transition matrix is thus $2^{T+m} = 16$.

$$\underline{\underline{T}} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0625 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.9375 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.1875 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.8125 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.1875 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.1094 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.8125 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.8906 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0312 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.9688 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.75 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0.25 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \tag{4.13}$$

Each non-zero element in the transition matrix corresponds to a probability function in the output space in the future-cast operator matrix. These are given explicitly in the Appendix A. Consider that the initial state is $\Gamma = 10$ i.e. the last 4 bits are $\{1001\}$ (obtained from running the game prior to the future-casting process). The initial probability state vector is the point (0,1) in the element of the vector $\underline{\varsigma}$ corresponding to time-horizon state 10. We can then generate the future-cast for given U (4.1).

Clearly the probability function in output space becomes smoother as U becomes larger and the number of successive convolutions increases, as highlighted by the probability distribution functions at $U = 15$ and $U = 25$ (4.2).

We note the non-Gaussian form of the probability distribution for the future-casts, emphasising the fact that such a future-cast approach is essential for understanding

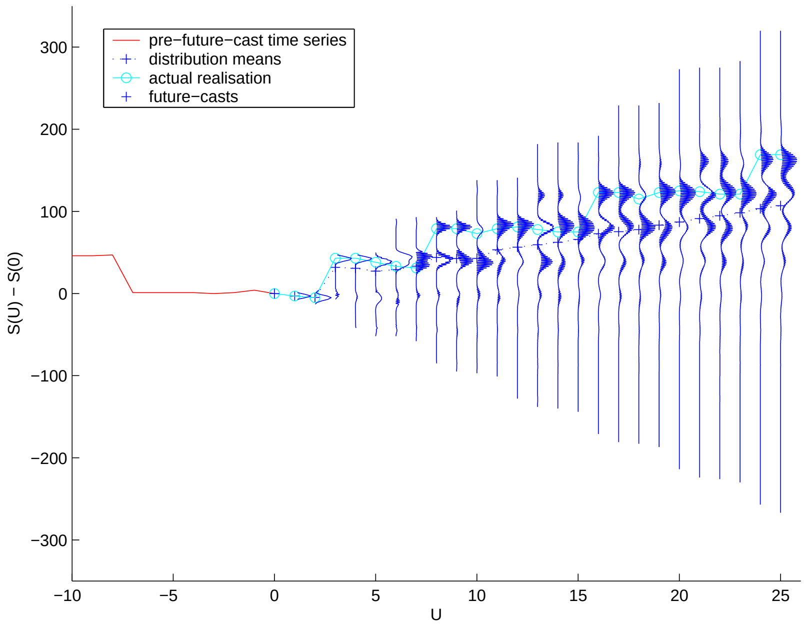

Figure 4.1: The (non-normalised) evolution of a future-cast for the given $\underline{\underline{\Omega}}$, game parameters and initial state Γ. The figure shows the last 10 timesteps prior to the future-cast, the means of the distributions within the future-cast itself, and also an actual realization of the game run forward in time.

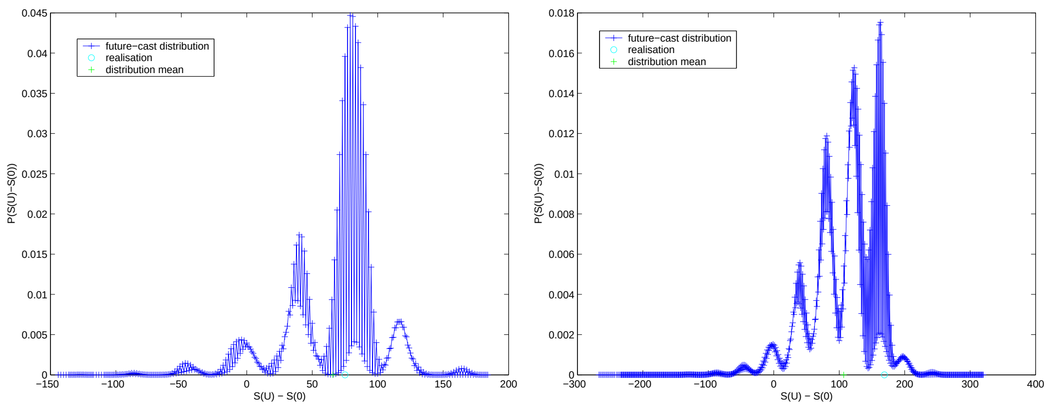

Figure 4.2: The probability distribution function at $U = 15, 25$ timesteps beyond the present state.

the system's evolution. An assumption of rapid diffusion toward a Gaussian distribution, and hence the future spread in paths increasing as the square-root of time, would clearly be unreliable.

Chapter 5

Evolution Management via Perturbations to Population's Composition

For less simple parameters, the matrix dimension required for the future-cast process become very large very quickly. To generate a future-cast appropriate to larger parameters e.g. $m = 3$, $T = 10$, it is still however possible to carry out the calculations numerically quite easily. As an example, we generate a random $\underline{\underline{\Omega}}$ (the form of which is given below) and initial time-horizon appropriate to these parameters. For visual representation reasons, the Reduced Strategy Space[8] is employed, the sampling method described in the appendix.This time-horizon is obtained by allowing the system to run prior to the future-cast. The other game parameters are as previously stated. The game is then instructed to run down every possible winning decision path exhaustively. The spread of output at each step along each path is then convolved with the next spread such that a future-cast is built up along each path. Fortunately, not all paths are realized at every timestep since the stochasticity in the winning-decision/state-space results from the condition $N_U \geq D_D$. The future-cast as a function of U and $S - S_{U=0}$, can thus be built up for a randomly chosen initial quenched disorder matrix (QDM) (5.1).

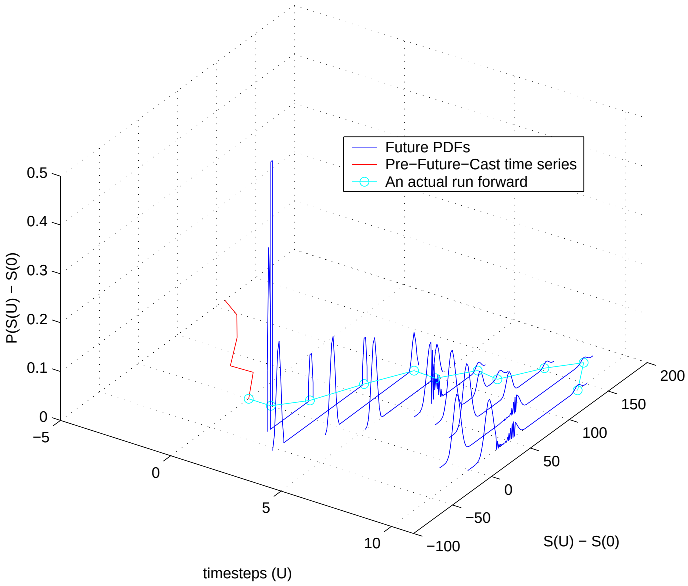

Figure 5.1: Evolution of $\Pi^U(S)$ for a typical quenched disorder matrix $\underline{\underline{\Omega}}$.

We now wish to consider the situation where it is required that the system should not behave in a certain manner. For example, it may be desirable that it avoid entering a certain regime characterised by a given value of $S(t)$. Specifically, we consider the case where there is a barrier in the output space that the game should avoid, as shown in Figure 5.2.

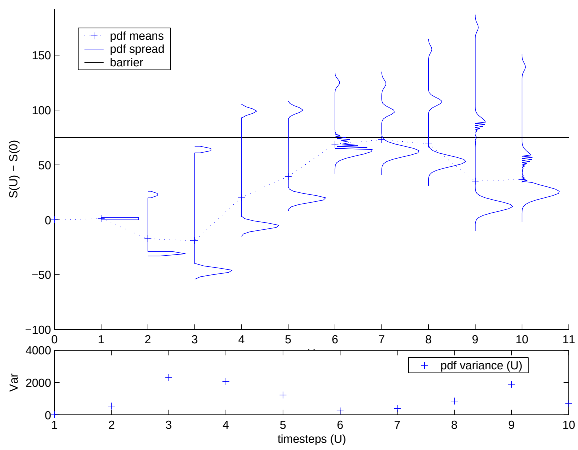

Figure 5.2: The evolution of the future-casts, and the barrier to be avoided. For simplicity the barrier is chosen to correspond to a fixed $S(t)$ value of 110, although there is no reason that it couldn't be made time-dependent. Superimposed on the (un-normalised) distributions, are the means of the future-casts, while their variances are shown below.

The evolution of the spread (i.e. standard deviation) of the distributions in time, confirms the non-Gaussian nature of the system's evolution – we note that this spread can even decrease with time[1]. In the knowledge that this barrier will be breached by this system, we therefore perturb the quenched disorder at $U = 0$. This perturbation corresponds in physical terms to an adjustment of the composition of the agent population. This could be achieved by 're-wiring' or 'reprogramming' individual agents in a situation in which the agents were accessible objects, or introducing some form of communication channel, or even a more 'evolutionary' approach whereby a small subset of species are removed from the population and a new subset added in to replace them. Interestingly we note that this 'evolutionary' mechanism need neither be completely deterministic (i.e. knowing exactly how the form of the QDM changes) nor completely random (i.e. a random perturbation to the QDM). In this sense, it seems tantalizingly close to some modern ideas of biological evolution, whereby there is some purpose mixed with some randomness. We leave a fuller discussion of this point to the future.

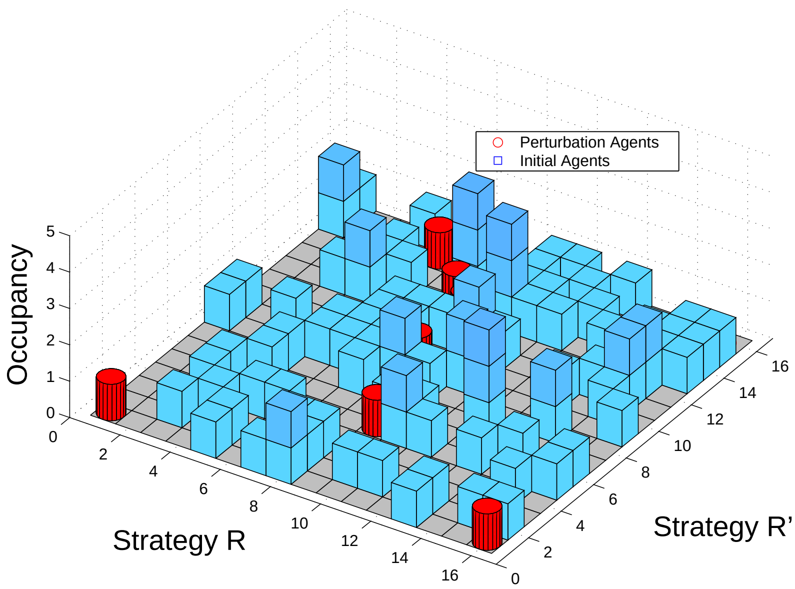

Figure 5.3: The initial and resulting quenched disorder matrices (QDM), shown in schematic form. The x-y axes are the strategy labels for the two strategies. The absence of a symbol denotes an empty bin (i.e. no agent holding that particular pair of strategies).

[1]This feature can be understood by appreciating the multi-peaked nature of the distributions in question. The peaks correspond to differing paths travelled in the future-cast, the final distribution being a superposition of these. If these individual path distributions mean-revert, the spread of the actual future-cast can decrease over short time-scales.

Figure 5.4 shows the impact of this relatively minor microscopic perturbation on the future-cast and global output of the system. In particular, the system has been steered away from the potentially harmful barrier into 'safer' territory.

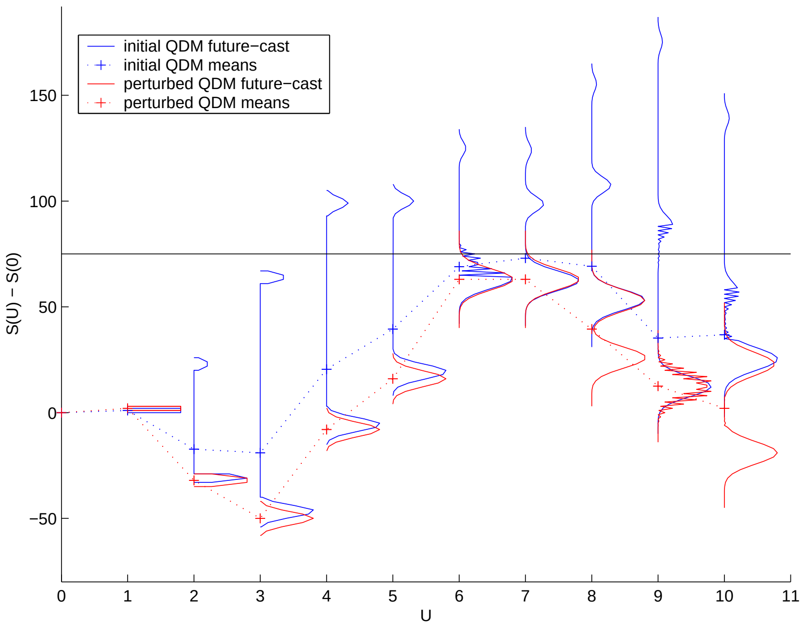

Figure 5.4: The evolution as a result of the microscopic perturbation to the population's composition (i.e. the QDM).

This set of outputs is specific to the initial state of the system. More typically, we may not know this initial state. Fortunately, we can make use of the characteristic future-casts to make some kind of quantitative assessment of the robustness of the quenched disorder perturbation in avoiding the barrier, since this procedure provides a picture of the range of possible future-case scenarios.

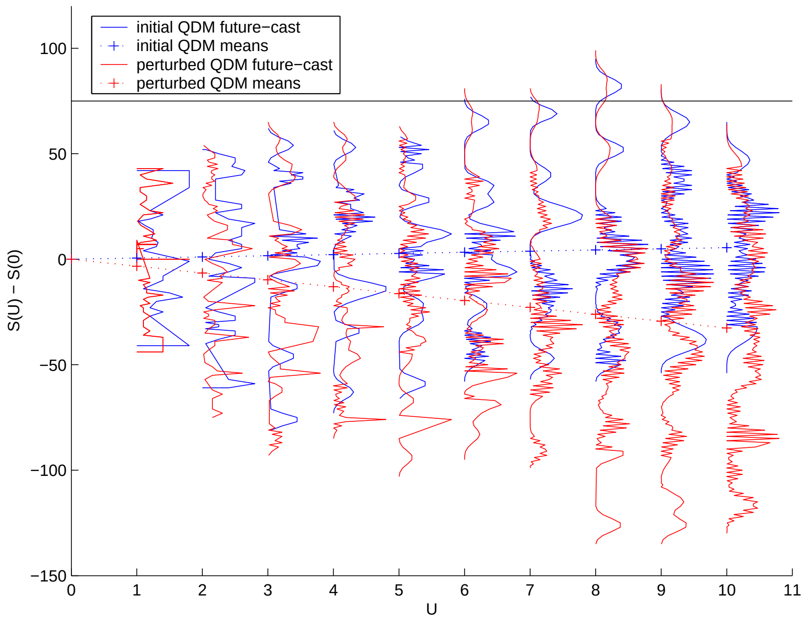

Figure 5.5: The characteristic evolution of the initial and perturbed QDMs.

This evolution of the characteristic future-casts, for both the initial and perturbed quenched disorder matrices, is shown in Figure 5.5 A quantitative evaluation of the robustness of this barrier avoidance could then be calculated using traditional techniques of risk analysis, based on knowledge of the distribution functions and/or their low-order moments.

Chapter 6

Concluding Remarks and Extensions

We have presented an analytical formalism for the calculation of the probabilities of outputs from the B-A-R system at a number of timesteps beyond the present state. The construction of the (homogeneous) future-cast operator matrix allows the evolution of the systems output, and other macroscopic quantities of the system, to be studied without the need to follow the microscopic details of each agent or species. We have concentrated on single realizations of the quenched disorder matrix, since this is appropriate to the behavior and design of a particular realization of a system in practice. An example could be a financial market model based on the B-A-R system whose derivatives could be analysed quantitatively using expectation values generated with the future-casts. We have also shown that through the normalised eigenvector solution of the Markov Chain transition matrix, we can use the future-cast operator matrix to generate a characteristic probability function for a given game over a given time period. The formalism is general to any time-horizon game and could, for example, be used to analyse systems (games) where a level of communication between the agents is permitted, or even linked systems (i.e. linked games or 'markets'). In the context of linked systems, it will then be interesting to pursue the question as to when adding one 'safe' complex system to another 'safe' complex system, results in an 'unsafe' complex system. Or thinking more optimistically, when can we put together two or more 'unsafe' systems and get a 'safe' one?

Future work will focus on the 'reverse problem' of generating a specific system to behave in a desired fashion. As such, the effects of any perturbation to the system's heterogeneity could be pre-engineered. A future application might include the global control problem with discrete actuating controllers[17].

Acknowledgements

DMDS is grateful to Royal Bank of Scotland for funding under an EPSRC Case studentship.

Appendix A

An Explicit Form of the Adjacency Matrix.

The elements of the adjacency matrix $\underline{\underline{\Upsilon}}$ are

$$\underline{\underline{T}} = \begin{pmatrix}
0&0&0&0&0&0&0&0&0&0&0&0&0&0&0&0\\
a_1\otimes&0&0&0&0&0&0&0&i_1\otimes&0&0&0&0&0&0&0\\
0&b_0\otimes&0&0&0&0&0&0&0&j_0\otimes&0&0&0&0&0&0\\
0&b_1\otimes&0&0&0&0&0&0&0&j_1\otimes&0&0&0&0&0&0\\
0&0&c_0\otimes&0&0&0&0&0&0&0&k_0\otimes&0&0&0&0&0\\
0&0&0&0&0&0&0&0&0&0&0&0&0&0&0&0\\
0&0&0&d_0\otimes&0&0&0&0&0&0&0&l_0\otimes&0&0&0&0\\
0&0&0&d_1\otimes&0&0&0&0&0&0&0&0&0&0&0&0\\
0&0&0&0&e_0\otimes&0&0&0&0&0&0&0&m_0\otimes&0&0&0\\
0&0&0&0&e_1\otimes&0&0&0&0&0&0&0&m_1\otimes&0&0&0\\
0&0&0&0&0&0&0&0&0&0&0&0&0&n_0\otimes&0&0\\
0&0&0&0&0&f_1\otimes&0&0&0&0&0&0&0&n_1\otimes&0&0\\
0&0&0&0&0&0&g_0\otimes&0&0&0&0&0&0&0&p_0\otimes&0\\
0&0&0&0&0&0&g_1\otimes&0&0&0&0&0&0&0&p_1\otimes&0\\
0&0&0&0&0&0&0&h_0\otimes&0&0&0&0&0&0&0&q_0\otimes\\
0&0&0&0&0&0&0&0&0&0&0&0&0&0&0&0
\end{pmatrix}. \tag{A.1}$$

and the functions are:

	Δ S	$P(\Delta S)$
a_1	-79	1
b_0	0	0.5
b_1	0	0.5
c_0	3	1
d_0	0	0.125
	2	0.0625
d_1	-6	0.0625
	-4	0.25
	-2	0.375
	0	0.125
e_0	0	0.125
	2	0.0625
e_1	-6	0.0625
	-4	0.25
	-2	0.375
	0	0.125
f_1	-44	1
g_0	0	0.125
	2	0.5
g_1	0	0.125
h_0	48	1

	Δ S	$P(\Delta S)$
i_1	-46	1
j_0	1	0.0625
j_1	-7	0.0625
	-5	0.25
	-3	0.375
	-1	0.25
k_0	44	1
l_0	4	0.125
	6	0.375
	8	0.375
	10	0.125
m_0	1	0.09375
	3	0.015625
m_1	-9	0.015625
	-7	0.09375
	-5	0.234375
	-3	0.3125
	-1	0.234375
n_0	0	0.03125
n_1	-8	0.0625
	-6	0.25
	-4	0.375
	-2	0.25
	-0	0.03125
p_0	0	0.5
p_1	0	0.5
q_0	78	1

Appendix B

Macroscopic Properties of the System.

We consider games of various memorysize whose other parameters are

Number of agents	N_{tot}	$(2^{m+1})^2$
Strategies per agent	s	2
Resource level	ϕ	0.5
Time horizon	T	10
Threshold	τ	0.0

and such that the QDM in the reduced strategy space is of uniform occupancy.

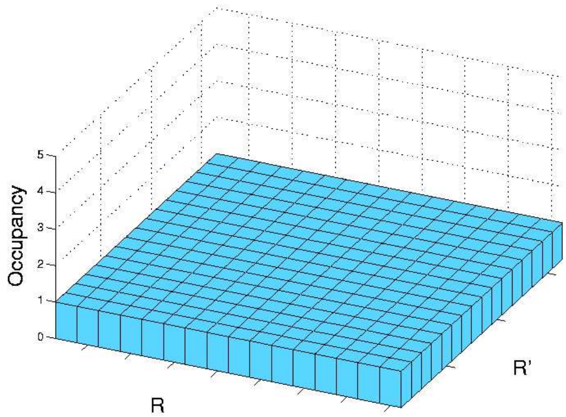

Figure B.1: The flat QDM in the reduced strategy space for a game of given m.

Using the future-casting framework, we generate the characteristic distributions of these games for a number of timesteps into the future, U. From these distributions, the macroscopic properties can be attained. Overleaf, the volatitities of these distributions is shown.

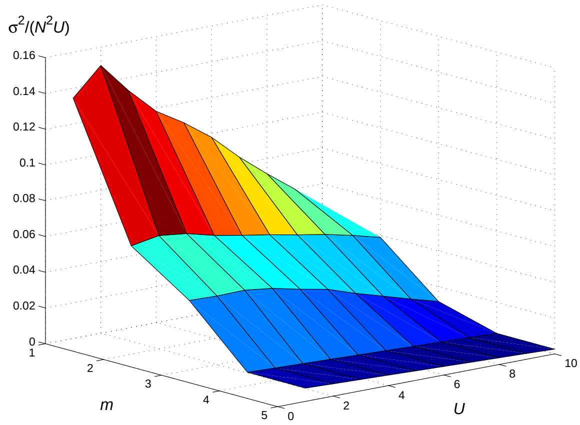

Figure B.2: Volatilities of the characteristic future-casts U timesteps into the future for the uniformly occupied strategy-space with varying memorysize.

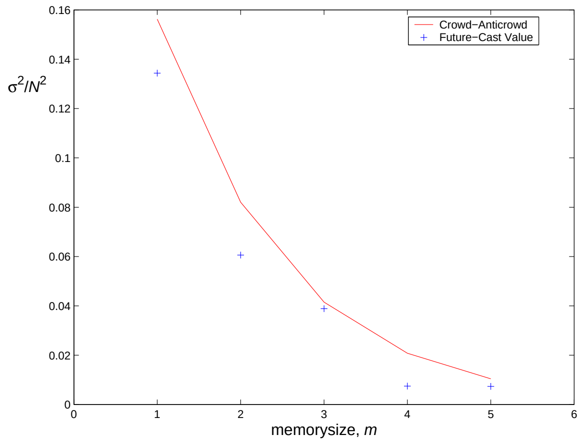

Figure B.3: The future-cast values (points) of volatility for $U = 1$ compared to the Crowd-Anticrowd model (line).

The results of B.2 compare well with those of the Crowd-Anticrowd theory of the minority game[9].

Bibliography

[1] D.H. Wolpert, K. Wheeler and K. Tumer, Europhys. Lett., **49**(6) (2000).

[2] B. Arthur, Amer. Econ. Rev. **84**, 406 (1994); Science **284**, 107 (1999).

[3] D. Challet and Y.C. Zhang, Physica A **246**, 407 (1997).

[4] D. Challet and Y.C. Zhang, Physica A **256**, 514 (1998).

[5] D. Challet, M. Marsili and R. Zecchina, Phys. Rev. Lett. **82**, 2203 (1999).

[6] D. Challet, M. Marsili and R. Zecchina, Phys. Rev. Lett. **85**, 5008 (2000).

[7] See http://www.unifr.ch/econophysics for Minority Game literature.

[8] N.F. Johnson, M. Hart and P.M. Hui, Physica A **269**, 1 (1999).

[9] M. Hart, P. Jefferies, N.F. Johnson and P.M. Hui, Physica A **298**, 537 (2001).

[10] N.F. Johnson, P.M. Hui, Dafang Zheng, and M. Hart, J. Phys. A: Math. Gen. **32**, L427 (1999).

[11] M.L. Hart, P. Jefferies, N.F. Johnson and P.M. Hui, Phys. Rev. E **63**, 017102 (2001).

[12] P. Jefferies, M. Hart, N.F. Johnson, and P.M. Hui, J. Phys. A: Math. Gen. **33**, L409 (2000).

[13] P. Jefferies, M.L. Hart and N.F. Johnson, Phys. Rev. E **65**, 016105 (2002).

[14] M.L. Hart, P. Jefferies and N.F. Johnson, Physica A **311**, 275 (2002) .

[15] M.L. Hart, D. Lamper and N.F. Johnson, Physica A **316**, 649 (2002).

[16] N.F Johnson and P.M. Hui, cond-mat/0306516v1 (2003).

[17] S. Bieniawski and I.M. Kroo, AIAA Paper 2003-1941 (2003).