CHALMERS
UNIVERSITY OF TECHNOLOGY

The Minority Game: evolution of strategy scores

M.Sc. thesis in the EMCAS program

Álvaro Pérez Díaz

Department of Applied Physics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2015

Master's thesis 2015

The Minority Game: evolution of strategy scores

ALVARO PEREZ DIAZ

Department of Applied Physics
Chalmers University of Technology
Gothenburg, Sweden 2015

The Minority Game: evolution of strategy scores
ALVARO PEREZ DIAZ

Supervisor: MATS GRANATH, Department of Physics, University of Gothenburg.

Master's Thesis 2015
Department of Applied Physics
Chalmers University of Technology
SE-412 96 Gothenburg

Typeset in LaTeX
Gothenburg, Sweden 2015

The Minority Game: evolution of strategy scores
ALVARO PEREZ DIAZ
Department of Applied Physics
Chalmers University of Technology

Abstract

The Minority Game is an agent based model that simulates competition for a scarce resource, situations in which two options are available to the agents at every time step and the winner option is the minority one. It was originally developed as a model for financial markets, although it has been applied in different fields, like genetics and transportation problems (choose less frequented road, lane, etc). This model has been studied in detail in the last fifteen years, with more than a thousand papers published on the topic, covering a wide range of analytical techniques, improvements and modifications, and in the recent years, large integration with different market mechanisms that reproduce the stylized facts of real markets.

We will first explain the model in detail and state its most important features, such as the existence of a phase transition that divides the game in two possible different regimes. We are interested in the so-called *dilute regime*, and we will describe in detail its particularities, which inspired our own work: it presents two different kinds of agents with very different behaviours, all depending on the random initial conditions. We will focus on the analysis of the strategy scores, which are the key factor determining which category an agent lies in. We use a probability theory to devise an analytical model for the so-called *coin-toss* limit inside this regime, and a phenomenological model that explains the behaviour of the strategy scores in the whole regime.

In the last chapter, we will introduce a similar game in which we constrain these mentioned strategy scores, yielding simplified dynamics with similar outcomes: the dynamics are trapped in typically small cycles in the state space, different cycles being present and depending on the initial conditions of the game.

Acknowledgements

I would like to thank Mats Granath for all the help and invaluable discussions held during this six months, creating a very stimulant and nice atmosphere.

Álvaro Pérez Díaz, Gothenburg, June 2015

Contents

List of Figures

1
Introduction

Traditionally, models in economics assumed *perfect, logical and deductive rationality.* The economist Brian Arthur is one of the biggest detractors of this way of thinking, claiming a *long frustration with mainstream theory's insistence that economic agents are equipped with such rational minds as to know everything, understanding all with implicitly infinite information capacity. He wanted to return to the drawing board, have a fresh look at agents with* bounded rationality [2].

The concept of *bounded rationality* was introduced by Herbert Simon in 1981 [3], it is the idea that *in decision-making, rationality of individuals is limited by the information they have, the cognitive limitations of their minds, and the finite amount of time they have to make a decision*, and Arthur thought that this fitted better our own minds, rather than 'god-like' agents.

He also claims that *inductive reasoning*, in which the agents are aware of an incomplete part of reality, and they build up their decisions trying to generalize this information matches better our own situation rather than the omnipresent *deductive approach*, in which agents descend to their particular situation from a general theory. It is usually explained as inductive meaning moving from specific observations towards broader generalisations and theories, whereas deductive moves from a general theory to the specifics. Thus inductive reasoning is conceptually much closer to the way human beings take decisions, as we have limited cognitive abilities and cannot usually comprehend systems as a whole, at least systems with massive number of variables such as financial markets.

To illustrate his views, Arthur developed the *El Farol Bar* problem [1], which goes as follows: El Farol Bar is an Irish bar in Santa Fe that has live music on Thursdays, and in his own words: *there are 100 Irish music lovers but El Farol has only 60 seats. The show is enjoyable only when fewer than 60 people show up. What should people do?* To make the decision, the agents can use some historical information, like the number of attendants in the last several weeks, and use simple predictors to decide whether to attend or not the incoming week.

This very simply posed problem presents remarkable features: for example, there cannot exist an a priori best strategy, because then everybody will use it and it will not be useful, hence heterogeneity across agents arises. All this inspired the *Minority Game*, similar to the *El Farol bar* problem, but presenting a more rigorous framework and a more detailed formalism. It was developed by Damien Challet, Yi-Cheng Zhang and Matteo Marsili in 1997, and since then, more than a thousand papers have been published about it and it has inspired lots of different games and models.

It will be described in detail in the next chapter, together with some of the early key results. In the following chapters our work on the evolution of strategy scores and a modification of the original game are detailed.

2
The Minority Game (MG)

The El Farol bar problem does not specify very rigorously what the *predictors* are, this is the strategies agents use to choose whether they attend or not. In the original formalism, a predictor *is a rule which 'predicts' the attendance of the next week, given the information about the attendances of the past weeks. Each agent has more than one predictor, he ranks them according to their performance and follows the recommendation of the best one* [2].

As mentioned in the previous chapter, the *Minority Game* developed by Challet, Zhang and Marsili in 1997, presents a precise framework with precise definitions. We will devote the rest of this chapter to go through basic definitions and main features of the MG.

2.1 Description

In the basic Minority Game, an odd number N of agents compete in successive rounds where they can choose between two options, say -1 or 1, wanting to be in the minority side each round. If they succeed they get a reward, otherwise they get punished. At the beginning of the game each agent is assigned a number of random strategies, which will govern the agent's behaviour. To choose what strategy to use each round, each is assigned a score based on how well it has performed so far, the one with leading score is used at a time step. Breaking ties between strategy scores can be done in different ways, but in the original game it is resolved randomly.

This can be put more precisely as follows: all the N agents have a memory M, meaning that they can remember the winning side of the last M rounds of the game, hence there are 2^M possible past histories. A particular strategy can be seen as a table with 2^M rows, one for each possible past history, and the corresponding predictions, which will dictate the side to choose in every different possible past.
Every agent gets a number S of strategies drawn at random from the strategy pool, which contains 2^{2^M} different strategies. The case $S = 1$ is not very interesting in the sense that agents hold only one strategy and cannot evolve or learn. It still presents chaotic behaviour extremely sensitive to the initial conditions. From now on we will consider $S = 2$, as the same overall qualitative behaviour is obtained using any $S \geq 2$. We assign a score to each of the agents, and to each of the strategies, U, which will be updated every time step, where different payoff schemes can be chosen. Note that we update the scores of *all* strategies at all time steps

Table 2.1: Strategy example, $M = 2$

History	Prediction
00	1
01	1
10	0
11	1

independent of whether they have been used. Let's call $\mu(t)$ the current past history at time t, and denote the prediction of the strategy s of the agent i to the past string μ as $a_{i,s}^{\mu(t)}$. We define the *attendance* A as the collective sum of actions from all the agents at a time step:

$$A(t) = \sum_{i=1}^{N} a_{i,s_i(t)}^{\mu(t)}$$

where $s_i(t)$ is the best stragey hold by agent i at time t:

$$s_i(t) = \arg\ \max_s U_{i,s(t)}$$

When every agent has made its choice, we sum all the bets for the two sides and we find the minority choice (recall that N is chosen to be odd, so there is always a minority side) and reward/punish agents and strategies accordingly. The basic score payoff system is detailed below for the two choices of labels for the possible agent actions:

$$U_{i,s}(t+1) = U_{i,s}(t) - \text{sign}[a_{i,s}^{\mu(t)} A(t)],$$

looking at the formula we see that it just updates the previous score by adding $+1$ if the strategy has predicted correctly this time step, and subtracts -1 if it has not. In other texts the *sign* function is removed so the scores depend on the actual size of the minority side, being the reward/punishment larger the smaller the minority side. Both possibilities yield the same overall behaviour.

We will keep an analogous score system for agents, based on their predictions been right or not. To initialize the game, one can just draw a initial history value randomly. More about this initial condition can be found in chapter 5.

We can see that the game is *inductive*, as every agent bases its decisions on the (probably partial) best choice they know, and *adaptive* as the agent's actions evolve with time and past history.

2.2 Major features

Now that we have precisely defined our game, it is time to look at the quantities of interest and its main properties: we have defined the *attendance* A, which sums up the collective behaviour of all the agents at each time step. The attendance will then have $N+1$ possible values, and due to the symmetry between the two actions in the game, we find that its average value will simply be

$$\langle A(t) \rangle = 0$$

which does not carry any significant information.

2.2.1 Phase transition

The situation is completely different when we look at its variance, denoted by σ^2, which was the first quantity of interest studied for the MG, altogether with the quantity $\alpha = 2^M/N$, discovered in 1999 by Savit *et al.* [5] and by Challet *et al.* [11]. It was shown that the behaviour of the game does not depend on N or M independently, but only on their ratio given by α. When we look at the behaviour of the variance of the attendance with respect to α, which reflects one of the major features of the MG: the existence of a phase transition between different regimes, as shown below:

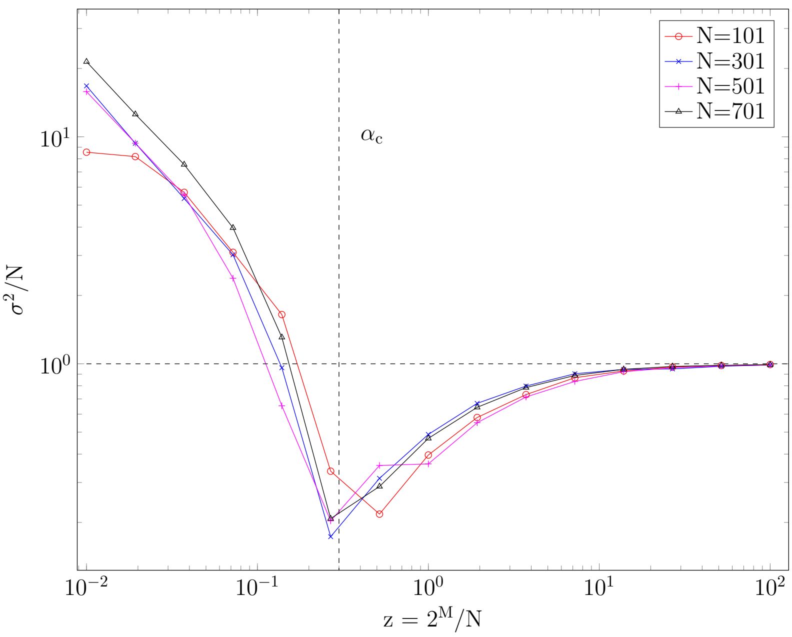

Figure 2.1: Phase transition with α as control parameter. Horizontal dashed line indicates coin-toss limit.

We can see the collapse of the three curves for different values of N and M. Three regimes are found:

- Small number of agents with respect to number of possible histories ($\alpha \gg 1$): called as well coin-toss limit, as the results seem random and similar to the ones obtained using a random coin-toss. This is produced due to over-fitting the past fluctuations of the attendance, as the information the agents receive at each time step is too complex to properly interpret it. This regime is commonly referred to as *dilute phase*, as there are few agents compared to the number of possible strategies.

- Increasing the number of agents ($\alpha \sim 1$): the fluctuations in the attendance are greatly reduced, showing that all the agents manage to coordinate themselves. As we decrease α from the coin-toss limit, σ decreases till it reaches a minimum around 0.34 for $S = 2$. //Challet and Marsilli 1999//.

- Increasing even more the number of agents ($\alpha \lesssim 0.1$): we see that the more we decrease α the larger the variance becomes, indicating large fluctuations in the attendance. This is related to the so-called *crowd effects*: as there are many agents sharing a small number of possible strategies (as M will be small in comparison), synchronized actions become usual. For this reason this regime is also known as the *crowded phase*.

Concerning the coin-toss limit, it is modelled as a binomial random variable $Bin(n, p)$, with $n = N$ number of agents and $p = 1/2$. Therefore the variance is found to be $\sigma^2 = 1/2 \cdot (1 - 1/2) \cdot 4 \cdot N$, thus $\sigma^2/N = 1$.

2.2.2 Predictability

In 1999, the same team led by Savit [5], discovered a new quantity featuring the same phase transition, the so-called *predictability*: firstly, consider $\mathbb{P}(1|\mu)$, the conditional probability of finding '1' as winning side given a previous history μ. This probabiliity is obtained by running a long simulation, storing the time series of past histories and respective outcomes, there we obtain a histogram with 2^M bins (possible past histories) ranging between 0 and 1. Unexpected behaviour is found here as well:

- $\alpha < \alpha_c$: on the left-hand side of the phase transition, all histograms $\mathbb{P}(1|\mu)$ look like shown below:

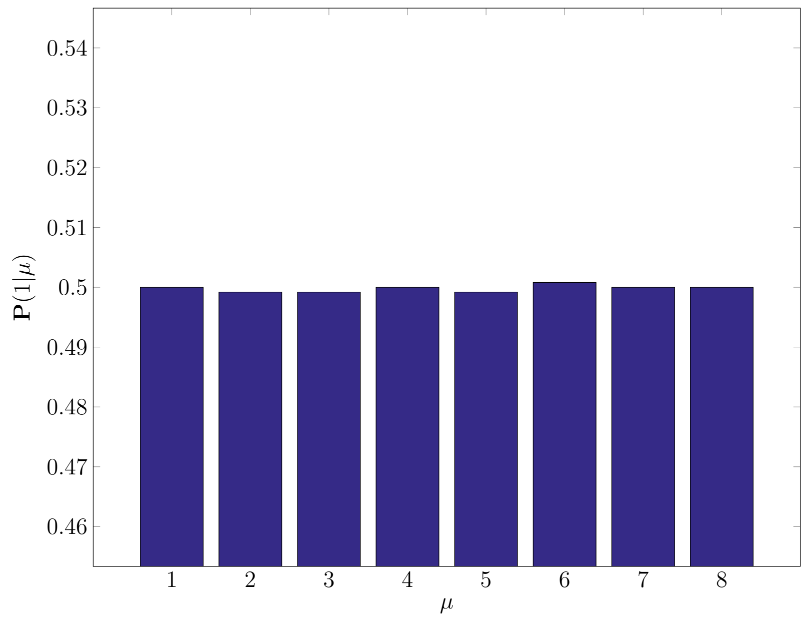

Figure 2.2: $\mathbb{P}(1|\mu)$ histogram for a realization of the MG with parameters $N = 51$ and $M = 3$, $\alpha = 0.1569$.

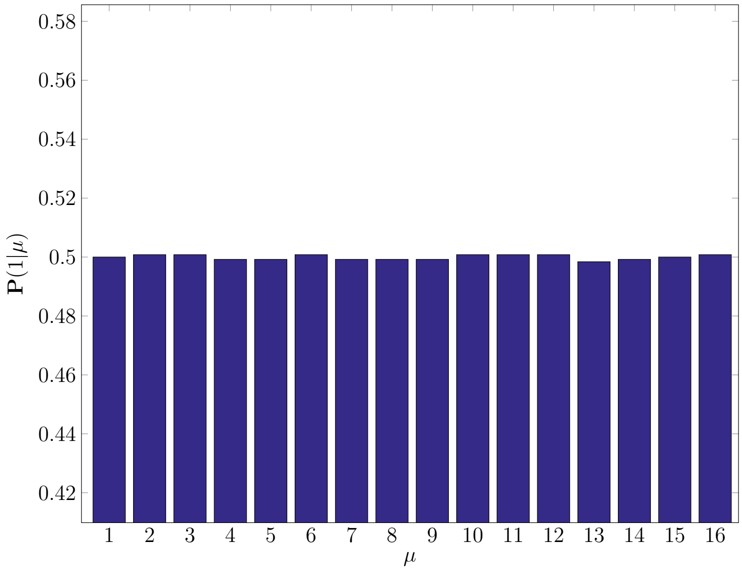

Figure 2.3: $\mathbb{P}(1|\mu)$ histogram for a realization of the MG with parameters $N = 51$ and $M = 4$, $\alpha = 0.3137$.

Both histograms were generated using 10^4 time steps, and we can see they are flat at probability 0.5, meaning there game is not predictable in this regime, also called the *unpredictable phase*. So for $\alpha < \alpha_c$ there is no extractable information from any history string of length M (the actual memory of the agents), as the probabilities of getting either action are equal.

- $\alpha > \alpha_c$: however, on the right-hand side of the critical point, the situation is entirely different: We can see that now each history string (of length M)

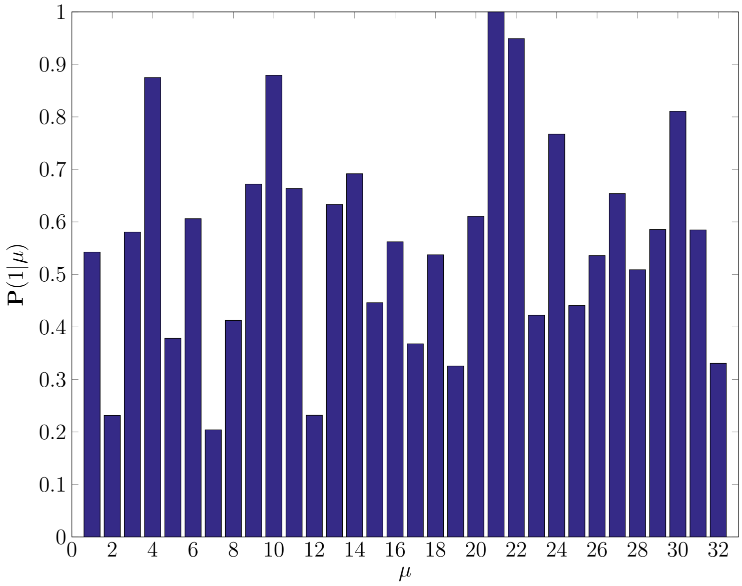

Figure 2.4: $\mathbb{P}(1|\mu)$ histogram for a realization of the MG with parameters $N = 51$ and $M = 4$, $\alpha = 0.3137$.

carries information about the best action for each past history, so we call this regime the *predictable phase*.

It was also found that if one looks at history strings with one more bit, this is, length $m + 1$, we will find predictitive information in both sides of the phase transition.

2.3 The Prior MG

The original formulation of the MG, as it has been defined in this chapter, can be seen as the sum of two processes: one *deterministic* and one *stochastic*. The stochastic term corresponds to randomly breaking ties between strategy scores. The deterministic term accounts for all the remaining dynamics, arising from the initial distribution of strategies to the agents, and the initial history.

In 2013 a modified game is introduced [9], without stochasticity, named *Prior MG* (PMG). In this game, instead of breaking ties randomly, each agent has a *preferred* strategy which they will use in case of tie.

Both games, MG and PMG, have the same overall qualitative behaviour, thus being interchangeable. We can use the simplified dynamics of the PMG to study properties of the MG. From now on in this report we will use exclusively the PMG, although we may use MG to refer to it.

3

The evolution of strategy scores

As mentioned before, the behaviour of the MG is very different in both sides of the phase transition, and this will affect greatly the dynamics of the strategy scores. On the left hand side, overall dynamics are periodic, including game states and the time series of strategy scores, which are bounded. When one jumps to the right hand side, however, things change as we find diverging strategy scores and not obvious periodic behaviour.

3.1 Divergence of strategy scores

When the MG is developing, each strategy has a virtual score which gets an update of ± 1 every time step. Theoretically, if a particular strategy works very good on this realisation of the MG, its score could just keep accumulating $+1$ updates and be always increasing towards ∞. Similarly for poorly suited strategies, towards $-\infty$.

To check the evolution of the strategy scores, we will run several realisations of the MG for different values of α, where we will check the maximum difference between strategy scores at the last time step. The results are shown below:

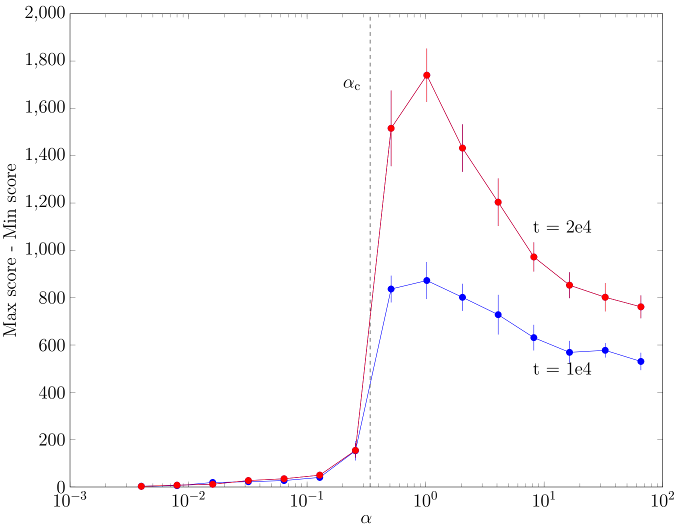

Figure 3.1: Max(score) - min(score) at the last time step for each value of α. Averaged over 10 different games with $t = 1e4$ (blue) and $t = 2e4$ (red).

We can see that the phase transition also affects the spread in the strategy scores, being tightly bounded for $\alpha < \alpha_c$, but allowing diverging spread when $\alpha < \alpha_c$. Of course, if one runs longer runs, this values will increment after α_c, but the shape curve remains intact.

When looking at the time series of strategy scores for a couple of realisations of the game, one for each regime, one finds the following:

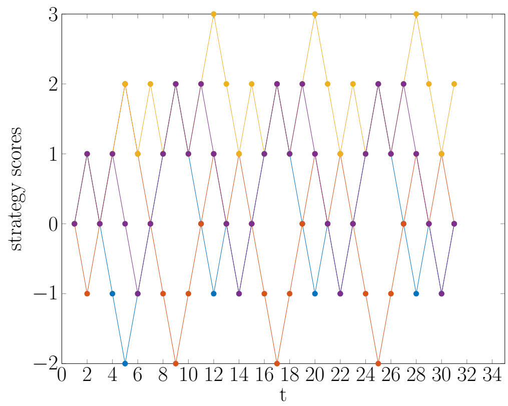

Figure 3.2: $\alpha = 0.0396$, two agents displayed, four strategies.

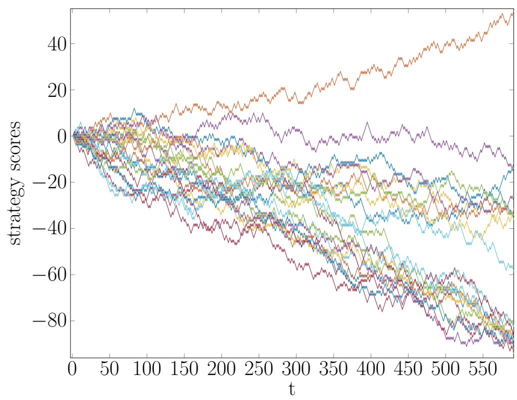

Figure 3.3: $\alpha = 0.7619$, 9 agents displayed, 18 strategies.

For the first figure, 3.2, we can see a completely periodic behaviour with period 8, after a very short equilibration time, and the strategy scores will keep repeating with this pattern. In the second figure, 3.3, we see that some strategy scores do grow nearly linearly towards infinity, while others stay around zero, and some unfortunate ones sink to minus infinity. This two examples show the stereotypical behaviour of all games with analogous values of α.

The period, if any, in the dilute phase becomes more difficult to asses: it greatly depends on the number of agents N and on α. Choosing small $\alpha \sim 1$ and small $N < 10$, one can easily find periodic behaviour with small period, but when one increases α or uses larger values for N the periods quickly scale to astronomical values which would need extremely long simulations to be found. There are other weaker periodic effects (quasi-periodic), see for example [9]: every consecutive times when we find the same past history, the minority side is very likely to switch, as the agents have learnt about the previous occasion. We will deal with this periodicity further in further chapters.

Concerning periodicity, the PMG behaves in a nicer way that the original MG: as the stochasticity is eliminated, we obtain periodic games (with small period) even in the dilute phase. In the MG, the presence of a small stochastic part, will generate quasi-periodic looking games, but not perfectly periodic as it is the case for the PMG as commented above.

3.2 Two types of agents

Taking any realisation of the PMG on the right hand side regime, we will see two kinds of agents, which have been named in the literature [11] (2008) as *frozen* and *fickle* agents:

- Frozen agents' strategies will have diverging scores, so that every time step (at least after some equilibration time) there will be a clear winning strategy which the agent will always use. This way, the behaviour of these agents becomes predictable and static.
- Conversely, fickle agents' strategy scores will be intertwined, in the sense that their scores will fluctuate around each other, causing the agent to actively switch strategies frequently.

An example can be found in the following picture, showing a realisation of the PMG with $N = 101$, $M = 7$ and $\alpha \approx 1.25$:

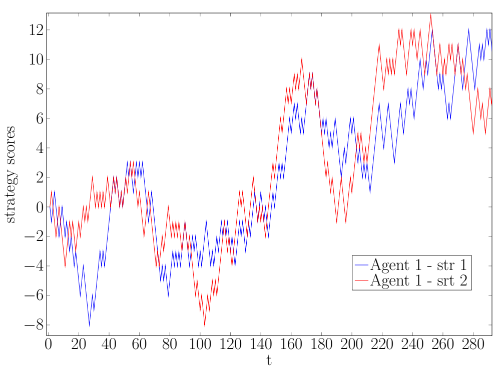

Figure 3.4: Two intertwined strategies from the same *fickle* agent.

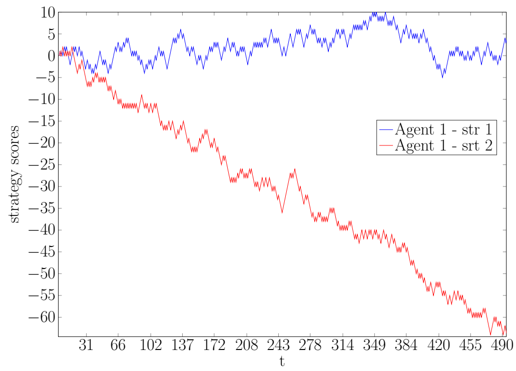

Figure 3.5: Two diverging strategies from the same *frozen* agent.

3.3 Intertwined strategies

The existence of these two types of agents is determined by the random distribution of strategies at the beginning of the game. When one of the strategies assigned to an agent performs way better than the other one, the agent will be frozen using that winning strategy and both strategy scores will diverge. This may not be the case for a different agent, which will switch between theirs.

When an agent is using a particular strategy, it is inducing a negative bias on it: it contributes to the outcome of the game this round, making it more likely to be in the majority group. Eventually if the strategy scores for both the strategies are close enough they will meet and possibilities of switching strategy will occur.

To get a better understanding of this phenomenon, we introduce the following experiment: we will run a normal simulation of the PMG for some time, find two fickle agents, and uncouple them from the game, which will run for some more time. With uncoupling we mean taking the two agents out of the game, in the sense that they are not counted towards the attendance anymore, so they do not participate in the minority and majority groups. But we still keep track of their strategies' scores. As they are effectively *not playing*, the negative bias they include in the strategies they use should be eliminated, and they strategy scores should diverge from each other, and follow their own paths depending on how good the strategies are suited for the current realisation of the game. One can afterwards include them back in the game to check this bias is acting again.

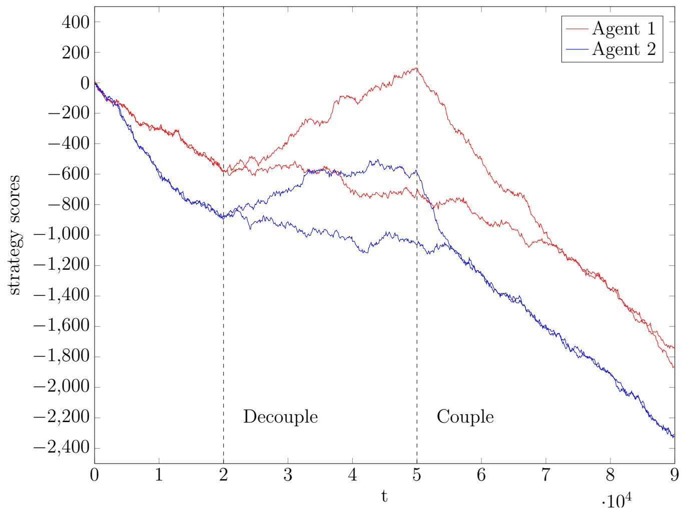

Figure 3.6: Typical behaviour when we decouple and couple fickle agents from the game.

We can see an example of this in figure (3.6), where the first and second dashed lines mark the decoupling and re-coupling times respectively.

Fickle agents will continue switching strategies forever, as their strategy scores will keep curling around each other. A similar phenomenon has been studied in [10], where the attractive behaviour between a strategy and its anticorrelated partner is detailed. It is shown that in the dilute phase we are looking at, *pairs of anticorrelated strategies would cycle around each other thus producing an ever-changing strategy rank vector*, where anticorrelated strategies s and s' verify $\sum_{\mu=1}^{2M} a_s^\mu a_{s'}^\mu = -2^M$.

Similar experiment can be performed for the frozen agents, one typical run of the game is displayed below in figure (3.7):
In this case, both agents use only one strategy, the one with higher score, which is the one located higher for every colour. We can see that the non-used strategies (the ones located lower for each colour) are not greatly affected by the decoupling: they are not in use at any time, However, we see that for both agents their preferred strategy grows much more when decoupled, as the negative bias introduced when they use it is removed. When we couple them again, we recover the original behaviour and leading strategies return to their previous score trends.

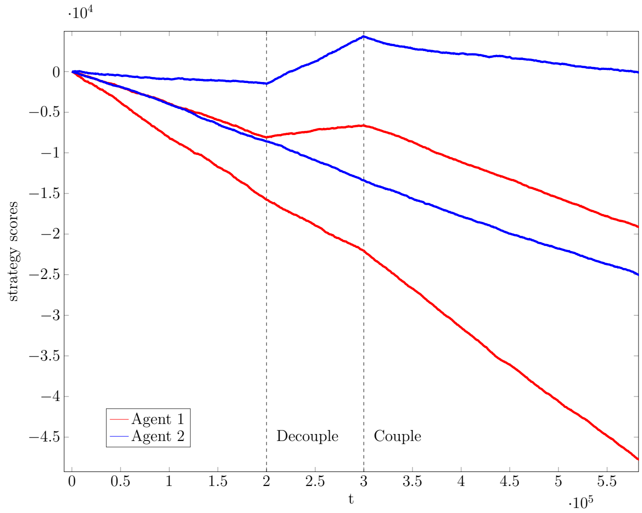

Figure 3.7: Typical behaviour when we decouple and couple frozen agents from the game.

All this observations led us to look at the distribution and behaviour of these pairs of strategy scores, which will be discussed in detail in the following chapter.

4

Analysis of the strategy score gaps

From now on, we shall refer to the score difference between an agent's two strategies as the *score gap* or just *gap*. Using the same notation as in the second chapter, we can write the time series of scores for the two strategies of agent i as $U_{i,1}(t)$ & $U_{i,2}(t)$. Thus the gap time series will be: $x_i(t) = U_{i,1}(t) - U_{i,2}(t)$.

Looking at the gap sign one will know which strategy the agent i is using: if $x_i(t) \geq 0$, the preferred strategy will be the first one, otherwise the second one will be selected.

Note as well that every time step the strategy scores are allowed to vary with ± 1 only, so the gap will vary with $-2, 0, 2$, which can be normalised to $-1, 0, 1$ just by taking $x_i(t) = (U_{i,1}(t) - U_{i,2}(t))/2$.

As we mentioned in the previous chapter, at the right hand side of the phase transition fickle agents will be forever switching strategies as the gap vector will keep switching from positive to negative values and vice versa. It then becomes interesting to investigate further on the distribution and behaviour of the gap time series.

4.1 The distribution of gap values

The first thing we need is to identify fickle and frozen agents, as the gap analysis only makes sense for the former. This can be accomplished by choosing a threshold value gap_{max} for the fickle agents' gap, then run a simulation checking whether the gap of the agents goes past the threshold, this will effectively distinguish between the two kinds of agents. To choose the threshold value one has to take into account the magnitude of the equilibration time, but still a large range of values works perfectly: if one chooses a too small threshold, we will be leaving fickle agents out, which will not be a problem, but we will get less data. For a too big threshold one will risk including frozen agents whose gap will just increment with time.

Once we have correctly identified the fickle agents we will continue the simulation storing the gaps for all of them at every time step. It is worth noticing here that we have arbitrarily assigned the labels 1 and 2 to the two strategies hold by each agent, so to eliminate this bias we will work all the time with the absolute value of the gap, making sure to double the zero gap, so that we do not double count in an inconsistent way.

We have performed extensive simulations for a variety of values of α, and in all cases the obtained distributions seem to follow, letting $x = |\text{gap}|$,

$$P_x \sim e^{-bx^c} \quad , \; x = 0, 1, 2, \ldots$$

the so-called *stretched exponential.*

Its parameters are a function of N and α, as we can see in the following plot for the exponent parameter c:

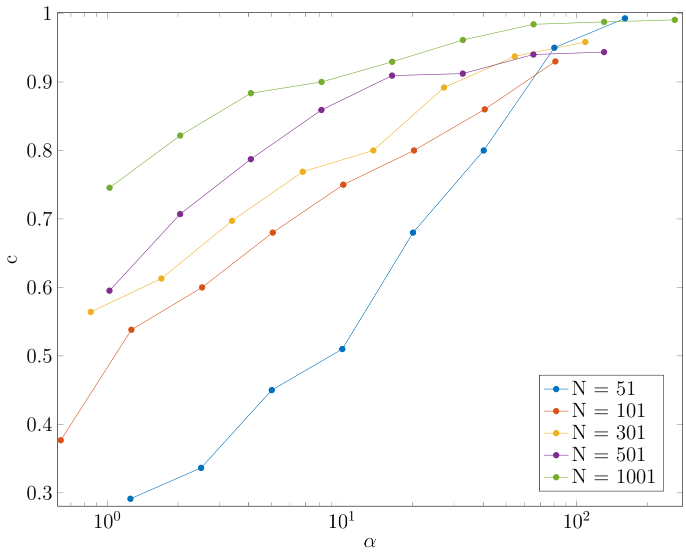

Figure 4.1: The parameter c of the stretched exponential as function of N and α.

We can see that for large N and large α, $c \to 1$, thus obtaining an exponential distribution. Two examples for low and large α are shown below:

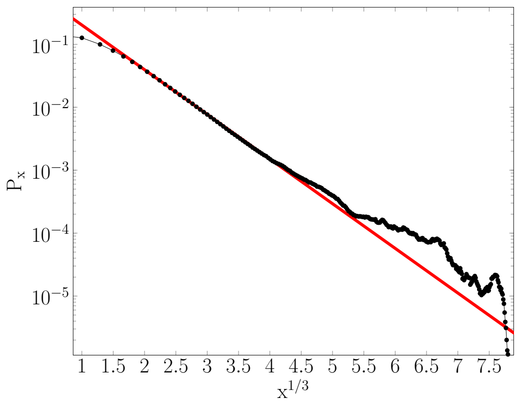

Figure 4.2: $N = 101$, $M = 7$, $\alpha \approx 1.27$, $c \approx 1/3$

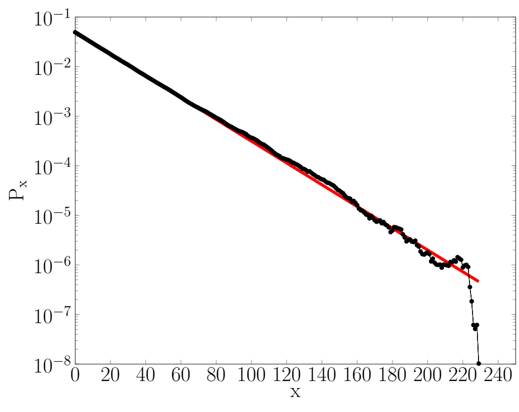

Figure 4.3: $N = 1001$, $M = 18$, $\alpha \approx 261$, $c \approx 1$

It is true, as we see in figure (4.2), that for small α the fit works very well for all but the very first values of x, which deviate themselves a bit from the stretched exponential. For larger α when the stretched exponential turns into exponential, the fit is great for all values of gaps x.

4.2 An upper bound for the maximum gap sizes based on time

In this chapter, we are studying the gap sizes in our simulations. It is possible then to introduce a simple estimate for an upper bound for the maximum gap size we expect to find, given the simulation time. The reasoning is as follows: given a time window τ, we expect to see a gap x if and only if $P_x \cdot \tau \geq 1$, thus

$$\tau \geq P_x^{-1} = \frac{1}{a} e^{bx^c} \Longrightarrow \log \tau \geq -\log a + bx^c \approx bx^c \Longrightarrow \boxed{x \leq \left(\frac{\log \tau}{b} \right)^{1/c}}$$

the timescale τ, in the periodicity sense, will be roughly the size of the state space (as the game is deterministic, it will repeat itself whenever it finds the same state, which cannot take longer than the size of the state space), which depends on the number of *fickle* agents N_s (the ones switching strategies, with dynamic behaviour), the maximum gap size x and the number of possible past histories:

$$\tau \approx 2^M (2x+1)^{N_s} \Longrightarrow \log \tau \approx N_s \log 2x$$

Hence, bringing together the two equations we obtain:

$$x \leq \left(\frac{N_s \log 2x}{b} \right)^{1/c} \Longrightarrow \boxed{N_s \log 2x - bx^c \geq 0}$$

This is a rough upper bound for the period time, as it is based on the game looping all over the possible state space, which includes all possible past histories and gap configurations for all the agents: the truth is that for most of the cases in the dilute phase, the game does not visit all the states evenly, so the estimate will be way bigger than reality, but still an upper bound.

4.3 The jump probability chain

We found experimentally that the the probability mass function for the size of the gaps fits a stretched exponential, $P_x \sim e^{-bx^c}$. To justify this experimental observation, we will describe in this section a phenomenological model that provides the stretched exponential behaviour.

We will look at the jumps between agent's strategy gaps: for every agent and time step of the game we will check the gap jumps (how the score difference between their two strategies varies that time step) and store them, so that we obtain the probabilities to jump to gap y from gap x, $\mathbb{P}(x \to y)$. As in a single time step the gap can only jump ± 1 or stay $+0$, we will effectively have three probabilities for each observed gap value x: (p_x^-, p_x^0, p_x^+), where $p_x^- + p_x^0 + p_x^+ = 1 \;\; \forall x$, see figure 4.4.

Experimentally we can run a simulation and obtain the probabilities (p_x^-, p_x^0, p_x^+) at each gap x for this game, the typical result looks like figure 4.5.

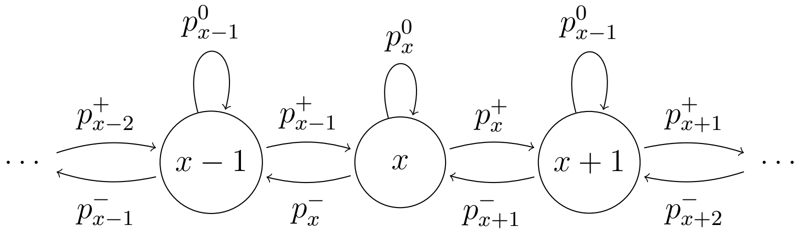

Figure 4.4: Jump probability chain, featuring possible gap values as nodes and the possible jumps with their probabilities.

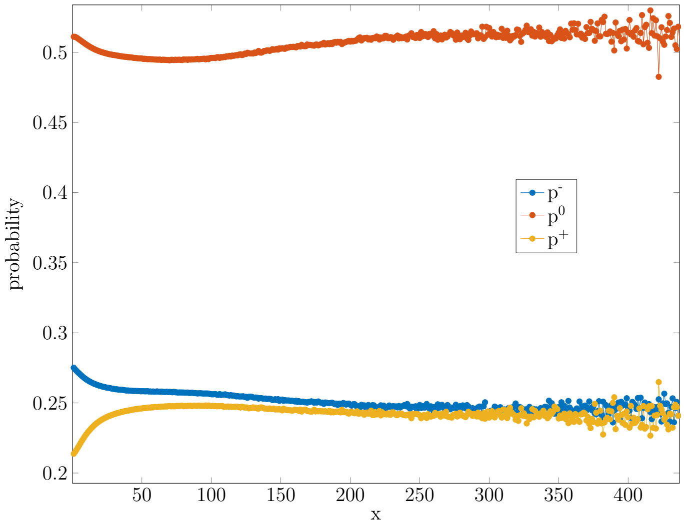

Figure 4.5: Set of probabilities (p^-_x, p^0_x, p^+_x) for each gap value.

This way we have build a chain with nodes $(0, \ldots, x_{\max})$, where each node has arrows pointing to the immediate right and left nodes and to itself. The next step will be to derive the master equation for this chain. Let $P_x(t)$ be the probability of occupation of node x at time t.

$$P_x(t+1) = p^0_x P_x(t) + p^-_{x+1} P_{x+1}(t) + p^+_{x-1} P_{x-1}(t)$$

We shall now solve the equation in order to find agreement with the observed experimental distribution $P_x(t)$. There is no experimental evidence of any time dependence on $P_x(t)$, so assume a stationary distribution and write $P_x(t) \equiv P_x$. If one has a look at the typical curves $p^-(x)$ and $p^+(x)$ we can see that they are nearly flat from gap values around 10, which induces the following assumptions: $p^+_{x-1} \approx p^+_x \equiv p^+$ and $p^-_{x+1} \approx p^-_x \equiv p^-$. Hence the master equation will look like:

$$P_x = p^0 P_x + p^- P_{x+1} + p^+ P_{x-1} = (1 - p^+ - p^-) P_x + p^- P_{x+1} + p^+ P_{x-1} \Longrightarrow$$
$$\Longrightarrow p^- \left(1 - \frac{P_{x+1}}{P_x}\right) = p^+ \left(\frac{P_{x-1}}{P_x} - 1\right)$$

This equation relates the probabilities $p^\pm$ to the gap size distribution P_x, so let's solve it. Trying the ansatz $P_x \sim \alpha^{-x} = e^{-x \log \alpha}$ we easily obtain:

$$\frac{p^-}{p^+} = \alpha \Longrightarrow \boxed{P_x \sim e^{-x \log\left(p^-/p^+\right)}} \tag{4.1}$$

and if one compares with the experimental expression $P_x \sim e^{-bx^c}$ one finds:

$$e^{-x \log\left(\frac{p^-}{p^+}\right)} \sim e^{-bx^c} \Longrightarrow \boxed{\log\left(p^-/p^+\right) \sim x^{c-1}}$$

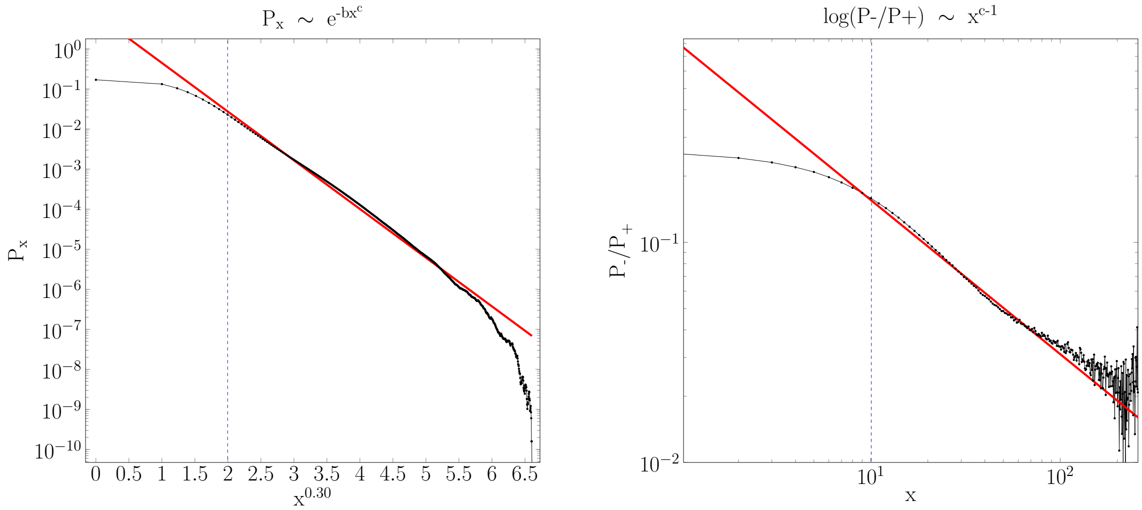

Figure 4.6: Fitting the model to a PMG with parameters $N = 101$, $M = 7$, $t = 10^9$.

In figure 4.6 one can see that our model fits the simulation data very well. The vertical dashed line represents the start of the gap values used to fit the stretched exponential to the probability mass function. As we commented earlier, when making the assumption of slow variation of $p_x^{+/-}$ with x, this would be reasonable for gap values greater than, approximately, ten, which is precisely the dashed line's position.

4.3.1 Random walk on the chain

As a experiment and confirmation of our chain analysis, once we have the jump probabilities for each gap (figure 4.5) we can perform a random walk on the gap sizes following these jump probabilities. This way we can obtain the probability of occupation for each gap size, which is equivalent to P_x. We show in the following picture an example run, where we can see very good agreement with the experimental probabilities even for small runs.

4.4 Analytical solution for the jump probability chain in the limit of large α

We remember from figure 2.1 that values of σ^2, for large α, approach $\sigma^2/N = 1$, as corresponding to a *coin-toss* limit. So in this regime, the non-used strategies will update their scores with ± 1 equally likely. For the strategies in use at a given time

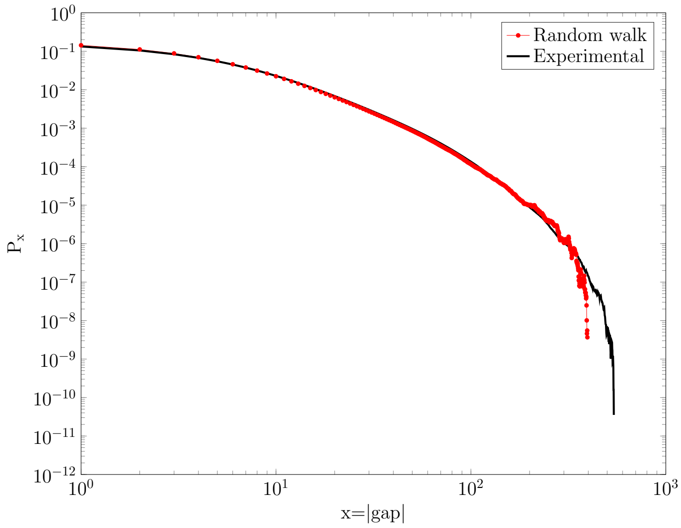

Figure 4.7: Random walk on the experimental chain. Parameters: $N = 101$, $M = 7$, 10^8 jumps.

step, we will have attendance:

$$A(t) = \sum_{i=1}^{N} a_i(t)$$

where $a_i(t)$ is the strategy used by agent i at time step t, and score updates:

$$U_i(t+1) = U_i(t) - a_i \cdot \text{sign}(A(t))$$

Without loss of generality, assume $a_i(t) = 1$ for a fixed value i and drop the explicit temporal dependence, then the attendance will look like:

$$A = 1 + \sum_{j \neq i} a_i$$

where the sum extends now over $N' = N - 1$ agents, being N' an even integer. However, as we work in the large N limit, $N \approx N'$ and we will drop the prime henceforth.
Let's look at the following probability:

$$P(A > 0) = P\left(\sum_{i=1}^{N} a_i \geq 0\right) = P(\lambda > 0) + P(\lambda = 0) \tag{4.2}$$

where $\lambda = \sum_{i=1}^{N} a_i$, the sum extending over all the agents but the fixed i-th one. As the strategies are randomly distributed, $P(\lambda > 0) = P(\lambda < 0)$, therefore

$$1 = 2 \cdot P(\lambda > 0) + P(\lambda = 0) \Longrightarrow P(\lambda > 0) = \frac{1}{2} - \frac{1}{2} \cdot P(\lambda = 0)$$

hence equation (4.2) turns into:

$$P(A > 0) = \frac{1}{2} + \frac{1}{2} \cdot P(\lambda = 0)$$

and similarly:

$$P(A<0)=\frac{1}{2}-\frac{1}{2}\cdot P(\lambda=0)$$

In terms of score jumps, keeping in mind that $a_i=1$:

$$P(A>0)=P(\Delta U_i=-1)\ \ \&\ \ P(A<0)=P(\Delta U_i=+1)$$

and defining again $x=(\Delta U_i^{+}-\Delta U_i^{-})/2$, the score gap between agent i's two strategies, where strategy U_i^{+} has the highest score (or either of them in case of tie), and $\Delta=x_t-x_{t-1}$ the gap variation at a given time step which can only be ±1, lets us write:

$$P(\Delta=1)=P(\Delta U^{+}=+1)\cdot P(\Delta U^{-}=-1)=\frac{1}{4}\left(1-P(\lambda=0)\right)$$

$$P(\Delta=-1)=P(\Delta U^{+}=-1)\cdot P(\Delta U^{-}=+1)=\frac{1}{4}\left(1+P(\lambda=0)\right)$$

$$P(\Delta=0)=P(\Delta U^{+}=-1)\cdot P(\Delta U^{-}=-1)+P(\Delta U^{+}=+1)\cdot P(\Delta U^{-}=+1)=\frac{1}{2}$$

With the same notation as in the previous section, we identify $p^{+}=P(\Delta=1)$, $p^{-}=P(\Delta=-1)$ and $p^{0}=P(\Delta=0)$.

Moreover, $P(\lambda=0)$ is binomially distributed, as it is just the probability that $N/2$ agents choose ±1 respectively, and one can write:

$$P(\lambda=0)=\binom{N}{N/2}\left(\frac{1}{2}\right)^{N}\approx_{(N\gg1)}\sqrt{\frac{2}{\pi N}}$$

where we have used Stirling's approximation for the factorial, valid for large N, $N\gg1$. Then we can approximate our previous formulae to:

$$P(\Delta=1)=p^{+}=\frac{1}{4}\left(1-\sqrt{\frac{2}{\pi N}}\right)\ \ \&\ \ P(\Delta=-1)=p^{-}=\frac{1}{4}\left(1+\sqrt{\frac{2}{\pi N}}\right)$$

and therefore:

$$\frac{p^{-}}{p^{+}}\approx\frac{1+\sqrt{\frac{2}{\pi N}}}{1-\sqrt{\frac{2}{\pi N}}}\approx1+\sqrt{\frac{8}{\pi N}}$$

Going back to equation (4.1), we obtain the following expression for the gap probability distribution:

$$P_x\sim e^{-x\cdot\log(p^{-}/p^{+})}\sim e^{-x\cdot\log(1+\sqrt{8/\pi N})}$$

and as $N\gg1$, we can Taylor expand the logarithm to finally obtain:

$$\boxed{P_x\sim e^{-x\sqrt{8/\pi N}}}\tag{4.3}$$

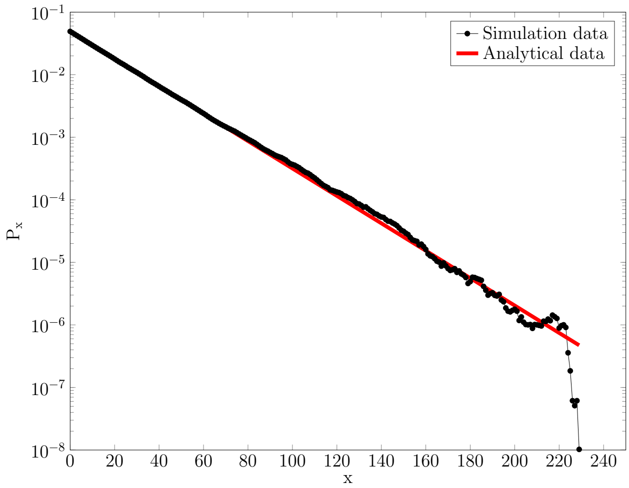

Figure 4.8: Analytical model for large α, PMG with parameters $N = 1001$, $M = 18$. The slope of the red line is $-\sqrt{8/\pi N}$.

One finds very good agreement between the analytical expression and simulation data in the right regime, large N and α. An example of a particular run of the game is displayed in the previous picture 4.8.

If one looks back at expression 4.3, we can rewrite it as $P_x \sim e^{-x/x_0}$, where

$$x_0 = \sqrt{\frac{N\pi}{8}}$$

is the characteristic length. We can then see that only for finite N the strategies are confined, in the $N \longrightarrow \infty$ we get a simple random walk (with $p^+ = p^-$) in the jump chain, yielding to diffusion behaviour: $\langle x^2 \rangle \sim t$ (time).

5

The Constrained Prior Minority Game (CPMG)

As studied in the previous chapter, frozen agents will be completely static as they do not switch strategies. Concerning the fickle agents, their score gap is enough information to know which strategy they are currently using. Moreover, when we use the *Prior* MG, we take out any kind of stochasticity from the simulation, being completely deterministic from the start. Although it still presents random *quenched disorder* corresponding to the randomly generated strategy tables.

These two features suggest the possibility of looking at the PMG as a *deterministic dynamical system* or Markov process, in which we can define a set of states and the deterministic transitions between them. A state is determined by the past history μ and the gap values for all the fickle agents $\{x_1, x_2, \ldots\}$. The frozen agents will always use the same strategy, so they do not count towards more states.

However, the possible gap values for the fickle agents are usually large, most of the cases of a few hundreds. This inspired us to introduce a slightly different game, which we named *Constrained Prior Minority Game* (CPMG). The difference with the usual MG is that we introduce a maximum gap value $x_{\max}$, which bounds the fickle agents' possible gap values. This maximum gap value can be chosen small compared to these typical few hundreds, thus obtaining a much smaller number of possible states and a greatly simplified behaviour in comparison with the PMG.

This way, the number of possible states will be given by $2^M(2x_{\max}+1)^{N_s}$, where N^s is the number of fickle agents, and we have used again the fact that gaps can only be even numbers, hence normalizing dividing by two so that we obtain successive gaps $\ldots, i-1, i, i+1, \ldots$. The 2^M comes from the number of possible past histories.

We can represent our system as a directed graph where the nodes are the different states and there is one arrow coming out from every node, which can be a self loop.

5.1 Justifying the use of the CPMG

When we bound the possible score gaps we are introducing a bias in the system, or equivalently, defining a new game. It is not obvious that both the classical and this newly defined games have the same behaviour, so that the CPMG is a good tool to study and analyse the PMG. There are several things we can check to guarantee this equivalence:

- The dependence on the initial history: we need a preliminary simulation to find the frozen agents, where we use a random initial history. Insensitivity to initial history will guarantee this procedure is reasonable. Also, getting similar score despite of this initial history will also show stability.

- The behaviour of the variance of the attendance, σ^2, one of the main features of the game. The new game exhibiting the same transition will guarantee the two games behave similarly.

We will go through each of these issues in detail now.

5.1.1 Initial history dependence

To check the effect of the initial history, one can distribute strategies to all players and then perform simulations of the same configuration of the game with each possible initial history, measuring whether this affects the agents being frozen or not. It is found that for all tested values of N: $101, 301, 501, 1001$; and α: from ≈ 1 to ≈ 120, being frozen is completely independent of the initial history, it is completely determined by the random assignment of strategies to all players. Then, it is safe to use a random history to characterise agent type.

As we can see in the following picture, agent success in the game depends mostly on their set of strategies, rather than the choice of initial history, as all the agents fluctuate very little for all histories:

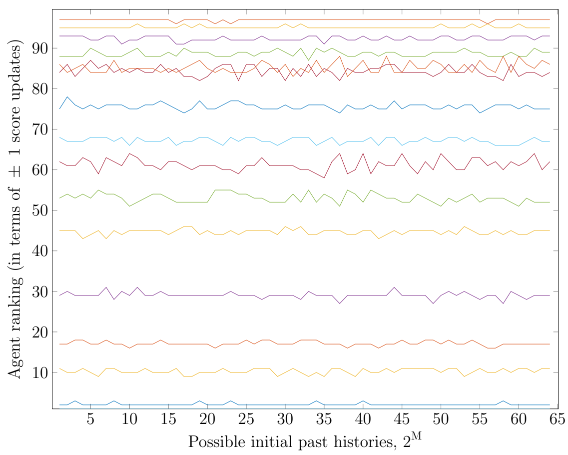

Figure 5.1: Agent ranking (from 1, best score, to N, worst) for each possible initial history. Each line is a different agent, a few agents have been picked and plotted to keep things clear.

5.1.2 The variance of the attendance

As explained in chapter (2.1), the variance of the attendance, σ^2, is one of the most important features of the MG. We would like to recover this two-phase structure for our game as well.

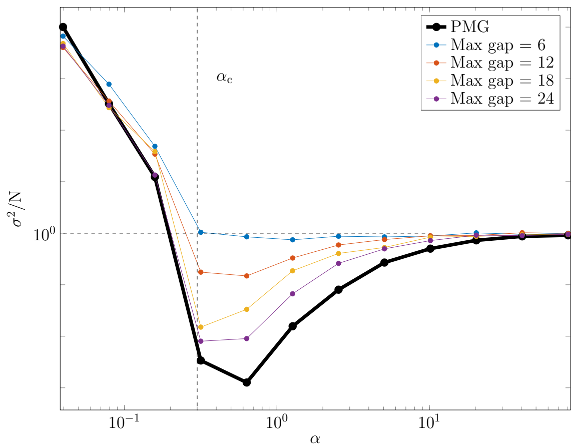

Figure 5.2: In thick black, the original PMG. Thin lines, CPMG with different $\max_{\text{gap}}$ sizes.

In figure (5.2) we can see that we recover the phase transition, the larger $x_{\max}$, the better resemblance compared to the original game. Clearly, the PMG is the $x_{\max} \longrightarrow \infty$ limit of the CPMG, but still the curves approach the original even for small $x_{\max}$.

5.2 Advantages of the CPMG

As we mentioned earlier, in the dilute phase, periods grow enormously big, making the game look absolutely aperiodic simulation-wise. However, when one constrains the possible gap values, the resulting game exhibits very small periods, reaching a stable periodic behaviour very quickly. We typically get a few connected components, two in the example plot below, and a few attracting short cycles, meaning from all states we reach a stable cycle with some small period ~ 10.

Depending on the initial state, the game will be quickly stabilised in one of this cycles. In the example below, figure 5.3, we find four different cycles, the game will be trapped in either of them depending on the initial history.

As we discussed in section 4.4, the PMG should exhibit much better periodic behaviour (for finite N) than the original game, which has a stochastic component. In the PMG we can find periodic behaviour (small period) for games in the dilute phase, keeping N very low, but as soon as we go to a bit larger games, the periods grow enormously. On the other hand, this CPMG keeps periods low, but it is true that its complexity and computation time grows fast with N and M.

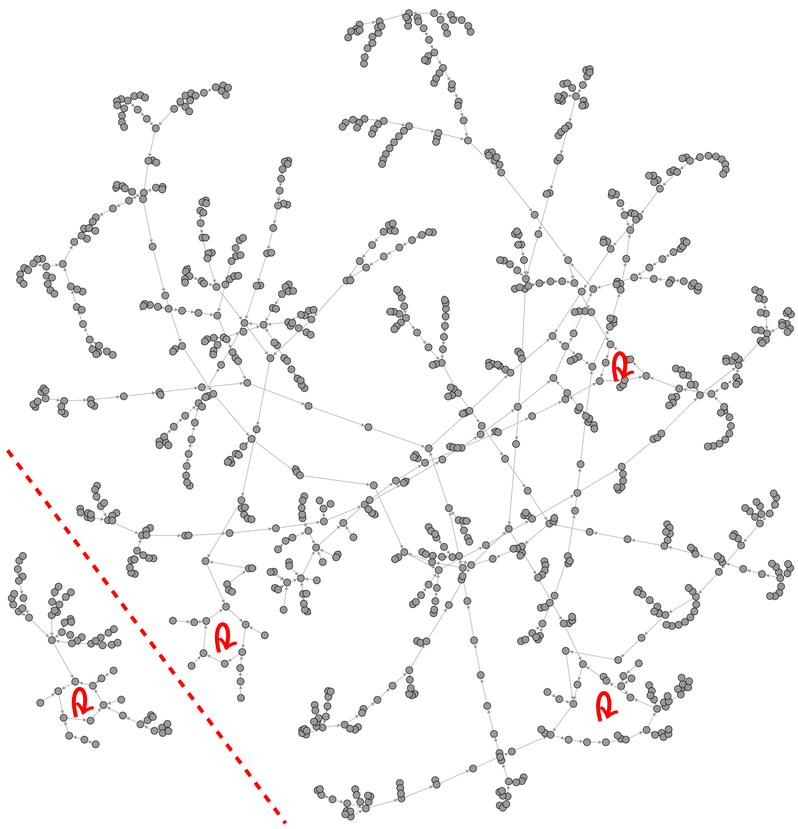

Figure 5.3: CPMG graph for parameters $N = 5$, $M = 3$, $\text{gap}_{\max} = 4$. Dashed red line separates the two connected components. Red loops show the attractive cycles.

6
Summary and conclusions

In the first chapter we presented a historical introduction to this kind of models inspired by economics, starting with *El Farol Bar* problem, which was the precursor of the *Minority Game*. It introduced the concepts of inductive reasoning and bounded rationality in a precise model which could be simulated mechanically.

The second chapter is devoted to the inner workings and important results concerning the MG: we started with the basic definitions to then introduce the phase transition and the predictability analysis. Also we described the modification to the original game with gives yield to the prior MG, getting rid of the small stochastic component of the early MG. This will be a key point as the two games are completely equivalent in terms of behaviour and will allow us to perform all the analytical analysis with greatly simplified statistical framework.

In the third chapter we start focusing on the evolution of the strategy scores, fundamental part of our work. Depending on the regime we are working at, we find very different overall behaviour which also translates in very different score behaviour: in the dilute phase we observe that the strategy scores diverge with time, rather than being bounded as happens in the crowded phase. This divides the agents into two types: frozen and fickle, depending on whether they only use one strategy continuously or they actively switching between their two. A key aspect of this fickle agents is the gap between their two strategies' scores, a big gap meaning they will use the leading strategy for a long time, whereas a small gap indicates faster switching.

In the fourth chapter we analyse the distribution and behaviour of this score gap: we performed extensive simulations to find that the distribution of the gap fits a *stretched exponential* distribution, which slowly turns into a regular exponential for large N and α, and a phenomenological model for this behaviour is provided solving the master equation of a random walk on jump chain. Moreover, an analytical model is detailed in the limit of large α, the *coin-toss* limit. We see very good agreement with the simulation data. Also an analytical upper bound to the period of the game is provided, explained in terms of basic probability and the game parameters.

In the fifth chapter, a modification of the MG is proposed: we will bound the gap between all the agents' strategy scores. This will create a finite state space of much smaller dimension compared to the original game, and as the game is deterministic we can represent the dynamics using a transfer matrix (equivalent to the adjacency matrix in graph theory). This gives yield to a much simplified behaviour, with small period behaviour. It can conveniently be presented as a graph, where the nodes are the possible states (in terms of past history and gap configuration for all the agents), with a single arrow coming out of every node indicating the next state.

An interesting problem for future work is to study similar analyitical models for the PMG in the case of the dilute phase but far from the coin-toss limit, where important correlations appear affecting greatly the behaviour of the agents. This analysis is much more complicated than the coin-toss model, and it will need extensive use of the *Central Limit Theorem*, to work statistically with the mentioned correlations.

Bibliography

[1] W. Brian Arthur. *Inductive Reasoning and Bounded Rationality: The El Farol Problem*. Amer. Econ. Review (Papers and Proceedings), 84, 406, 1994.

[2] D. Challet, M. Marsili, Y-C. Zhang. *Minority Games*, Oxford University Press, Oxford, UK, 2005.

[3] H. Simon. *The Sciences of the Artificial*, MIT Press. Cambridge, MA., 1981.

[4] D. Challet, Y.-C. Zhang. *Emergence of Cooperation and Organization in an Evolutionary Game*. Physica A: Statistical Mechanics and its Applications Volume 246, Issues 3–4, 1 December 1997, Pages 407–418.

[5] Robert Savit, Radu Manuca, Rick Riolo. *Adaptive Competition, Market Efficiency, and Phase Transitions*. Physical Review Letters, Volume 82, Number 10, March 1999.

[6] Robert Savit, Radu Manuca, Rick Riolo, Yi Li. *The Structure of Adaptive Competition in Minority Games*. Physica A: Statistical Mechanics and its applications, Volume 282, Issues 3-4, July 2000.

[7] D. Challet, Y.-C. Zhang. *On the Minority Game: Analytical and numerical studies*. Physica A: Statistical Mechanics and its applications, Volume 256, Pages 514-532, 1998.

[8] C.H. Yeung, Y.-C. Zhang. *Minority Games*. Encyclopedia of Complexity and Systems Science, Pages 5588-5604, Springer New York, 2009.

[9] G. Acosta, I. Caridi, S. Guala, J. Marenco. *The quasi-periodicity of the minority game revisited*. Physica A: Statistical Mechanics and its applications, Volume 392, Issue 19, Pages 4450-4465, 2013.

[10] P. Jefferies, M.L. Hart, N.F. Johnson. *Deterministic dynamics in the minority game*. Physical Review E - Statistical, Nonlinear, and Soft Matter Physics, Volume 65, Pages 1-8, March 2002.

[11] D. Challet, M. Marsili. *Phase transition and symmetry breaking in the Minoority Game*. Physical Review E - Statistical, Nonlinear, and Soft Matter Physics, 60, R6271(R), December 1999.

[12] D. Challet, M. Marsili. *Continuum time limit and stationary states of the Minority Game*. Physical Review E - Statistical, Nonlinear, and Soft Matter Physics, 64 056138, October 2001.

[13] T. Galla, G. Mosetti, Y-C. Zhang. *Anomalous fluctuations in Minority Games and related multi-agent models of financial markets.*. arXiv:physics/0608091 [physics.soc-ph], 2006.

[14] P. Jefferies, M.L. Hart, P.M. Hui, N.F. Johnson. *From market games to real-world markets*, The European Physical Journal B - Condensed Matter and Complex Systems, April 2001, Volume 20, Issue 4, pp 493-501.