The evolution of strategy scores in the Grand Canonical Minority Game

Master's Thesis in: Engineering Mathematics and Computational Science & Physics and Astronomy

Henrik Björkman & Andreas Olsson

Department of Applied Physics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2016

The evolution of strategy scores in the Grand Canonical Minority Game
HENRIK BJÖRKMAN & ANDREAS OLSSON

Supervisor and examiner: Mats Granath, Department of Physics, University of Gothenburg

Department of Applied Physics
Chalmers University of Technology
SE-412 96

The evolution of strategy scores in the Grand Canonical Minority Game
Henrik Björkman & Andreas Olsson
Department of Applied Physics
Chalmers University of Technology

Abstract

Understanding financial markets has always been a subject of interest for economists as well as physicists and mathematicians. In recent years there has been a lot of interest in agent-based models as a tool to understand market dynamics. Grand Canonical Minority Game (GCMG), a development of the Minority Game proposed in 1997, is an agent-based model where the agents, based on assigned strategies, choose one of two sides with the aim of choosing the minority side with the additional option to not participate (as opposed to the basic Minority Game). Such models have shown promise in regards to produce stylized facts from real markets such as heavy tailed price returns as well as volatility clustering around critical states.

In this thesis we first go through the basics of the Minority Game and show its features, then we move on to describe GCMG. In this game we have identified different kind of agents, which will be described in detail. The main part of the work has been to formulate a statistical model for the game which characterize the agents mean step size, with a step being how far an agent moves in a time step. From this we have been able to calculate fractions of these different kind of agents as well as the full score distribution of all agents. Our model qualitatively reproduces results from numerical simulations within a range of the reduced amount of speculators in the market ($n_s \leq 1$). This represents a game with a large strategy space compared to the amount of speculators.

Acknowledgements

We would like to thank Mats Granath for all the help and guidance in our work.

Henrik Björkman and Andreas Olsson, Gothenburg, December 2016

Contents

1 Introduction

The interest for agent based models of the financial market has grown a lot recently. Standard economic models have proved to be unable to anticipate financial crises and in light of this, alternative models are needed in order to get a better understanding of the financial market.

One model that has been used in the hope of gaining better understanding of the financial market is the Minority Game which takes its root in Brian Arthur's El Farol Bar Problem. The El Farol Bar Problem formulation is easy and can be summarized as [1]:
"N people wants to go to the El Farol Bar on a certain night each week where they have entertainment. However, if more than 60% of the people attend the bar, it will be too crowded and it would have been preferable to stay at home."
The people must make their choices, if they should attend the bar or not, independent of each other. One of the interesting features of this problem is that there can be no best strategy beforehand, because if there was everyone would use it and it would not have been the best. The people thus needs to use inductive reasoning. From this problem formulation Yi-Cheng Zhang and Damien Challet developed the Minority Game in 1997 [2], which we will discuss in detail in the next chapter.

We start this project by going through some basic features of the Minority Game. We then study the so called Grand Canonical Minority Game, which is a simple development of the Minority Game, to arrive at a statistical model for this particular game.

2 From El Farol to The Minority Game

In economics agents are thought of to have perfect rationality, this being that agents are equipped with a rational mind with infinite capacity to understand everything. Brian Arthur thought of the El Farol Bar Problem because he did not like this premise, instead he was interested in a problem where the rationality of the agents was bounded. The El Farol Bar Problem was thought as a way of thinking about inductive reasoning and its modeling [1][3]. Since the agents do not know each others choices they can not come up with a deductive solution to the problem. Hence the agents will have to utilize inductive reasoning. The expectations of the agents have to differ, they can not all be thinking the same way.

In the El Farol problem it was unspecified how the agents would predict whether or not to attend the bar. That is one of the things that become clearer in the Minority Game, as mentioned proposed by Challet and Zhang in 1997. The Minority Game is a game where agents compete over a scarce resource and reward those in minority while punish those in majority. The Minority Game has inspired a lot of research because of its interesting features. This chapter will describe the game and some its basic features.

2.1 Basics of the game

The Minority Game is an agent based model where an odd number, N, agents are choosing between two options, say 1 and -1 (think of it as buy and sell), with the aim of choosing the minority group. The agents make their decision based on a strategy. The strategy predicts the next winning bid given the last m number of previous winning bids. This information is called the history, μ. This suggests the size of the history space is $P = 2^m$. The agents that choose the minority group get rewarded whereas the agents in majority get punished. At the start of the game the agents are randomly assigned s strategies and in each time step they will use the strategy that has historically performed the best. If there is a tie between strategies it is either resolved by a coin-toss or alternatively one can predefine the choice of strategy in case of a tie.

What we have is an odd number of agents, N, who all have memory of the winning bid from the last m time steps. The agents strategies then predicts the winning bid for the next time step given any of these histories, a strategy when $m = 2$ can be seen in table 1.

History μ		Bid
-1	-1	1
-1	1	-1
1	-1	-1
1	1	-1

Table 1: An example of a strategy for $m = 2$.

This implies the complete strategy space is made up of a total of 2^{2^m} distinctly different strategies. Each agent is randomly assigned strategies from the strategy space in the initiation of the game. The number of different strategies, s, each agent are assigned is up for choice. However, as $s = 1$ renders a game where the agents cannot learn and adapt, one usually considers $s = 2$. The agents respective strategy is assigned a score, U, which will be updated after every time step. At the start of the game the score for each strategy is 0. An example of an updating scheme for the strategy score of agent i is

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

where $a_{i,s}^{\mu(t)}$ is the i'ths agent's choice, which is either 1 or -1, given the history $\mu(t)$ at time t and s signifies which strategy is being updated. $A(t)$ is called the attendance and is defined as

$$A(t) = \sum_{i=1}^{N} a_{i,s_i(t)}^{\mu(t)}, \tag{2}$$

where $s_i(t)$ is the i'ths agents best scored strategy at the particular time t. This implies the attendance is the combined bid of the agents. This particular update scheme is a sign payoff which gives the strategies that predicted the minority bid +1 and the other strategies -1. Note that every strategy gets a score update in every time step, including those which were not used. Another frequently used payoff scheme is the linear payoff which updates the score linearly to the attendance. We note that with either of these schemes the game is a negative sum game, that is the total score will always get lower, considering that more agents will be in majority than minority each time step.

2.2 Statistical properties

In the initiation of the game we randomize the agents strategies. This indicates that we can treat $a_i^{\mu(t)}$ as a random variable with equal probability of being 1 as

-1 and we have $\langle A(t)\rangle = 0$. Thus this quantity does not really tell us much of the game. The variance of $A(t)$ is a more interesting quantity since it gives us information of how effective the game (or market) is, how volatile it is if one will. The variance is calculated as $\sigma^2 = \langle A(t)^2\rangle - \langle A(t)\rangle^2 = \langle A(t)^2\rangle = \mathbb{E}[(\sum\limits_{i=1}^{N} a_{i,s_i(t)}^{\mu(t)})^2] = \sum_{i=1}^{N} \mathbb{E}[(a_{i,s_i(t)}^{\mu(t)})^2] = N$ whereas the volatility is defined as σ^2/N. Since the variance is calculated treating the agents choices as uniformly random variables, $\sigma^2 = N$ corresponds to a market where agents are acting randomly and if we have lower variance the market is thus more effective than agents picking randomly.

2.3 Controlling parameter

Previous studies of the game has shown that the parameter $\alpha = 2^m/N$ controls the game, in the sense that it controls the volatility [4]. In Figure 1, we see α plotted against σ^2/N. As we see, there is a phase transition between two regimes. The exact value of the phase transition has been found to be $\alpha_c \approx 0.34$ in the case of $s = 2$ [5]. When there is a high amount of agents compared to distinct strategies, we get crowd effects which leads to higher variance. Crowd effects simply being that since we have such a small strategy space the agents strategies will be fairly correlated leading to agents being grouped together.

As m grows, or N gets smaller, the variance lowers until the transition and then gradually rises until it reaches $\sigma^2/N = 1$ which is the so called coin-toss limit. In the coin-toss limit agents are essentially acting randomly on the market, having no correlation between their respective strategies. Around the phase transition we have a variance considerably lower which means the agents apparently makes informed decisions. One thing that has been studied quite extensively are models in which one let the agents learn along the way in order to get a market as effective as possible [6][7]. Learn in this context might be to be able to update their poor strategies from the best performing neighbour or something along that line. This could be interesting in applications since it would represent a stable market.

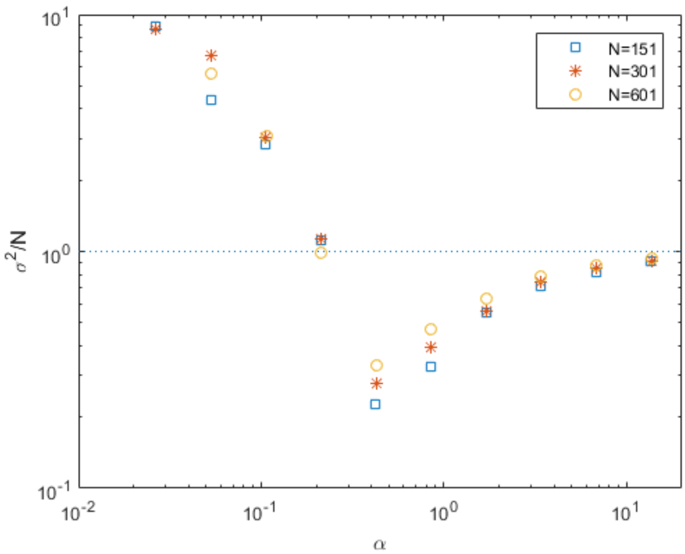

Figure 1: Here we can see the phase transition. For small α we have a large volatility whereas for larger values we get close to the coin-toss limit. Every point is averaged over 5 runs of 10^4 time steps.

2.4 Predictability

A reasonable question to ask oneself in light of these findings is how the agents around the phase transitions can make these informed choices. The answer to this questions is that there is a predictability in the game during this phase. We can see in Figures 2 and 3 the conditional probability of finding the agents to bid 1 given a history μ, $\mathbb{P}(1|\mu)$. We see that for small α we have no predictability whereas for higher values of α we have that the histories vary in probability and hence there is a predictability that the agents can use which results in lower variation of the attendance. Thus, the different phases in the game are usually referred to as the 'unpredictable phase' and the 'predictable phase'. Another way of looking at it is saying the phase over the vertical line in Figure 1 is the worse than random phase whereas under is the better then random phase. Drawing the parallel to the real market, the better than random phase is how we would prefer our market to be.

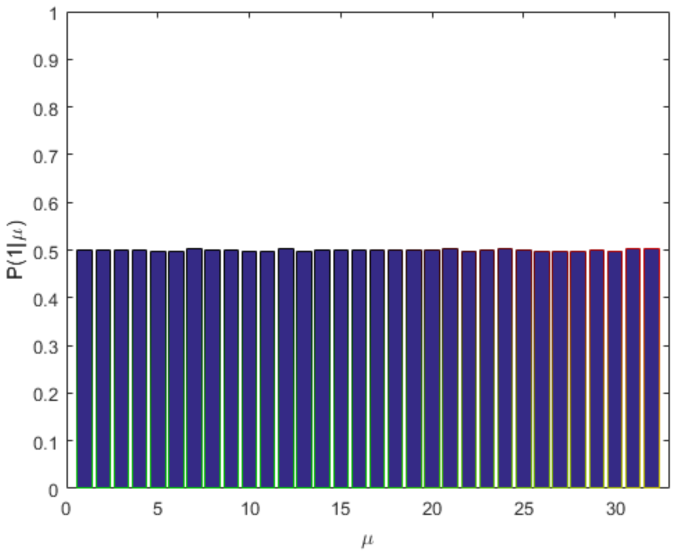

Figure 2: Histogram of conditional probability of choosing the bid 1 given a specific history for $\alpha = \frac{2^5}{501} < \alpha_c$.

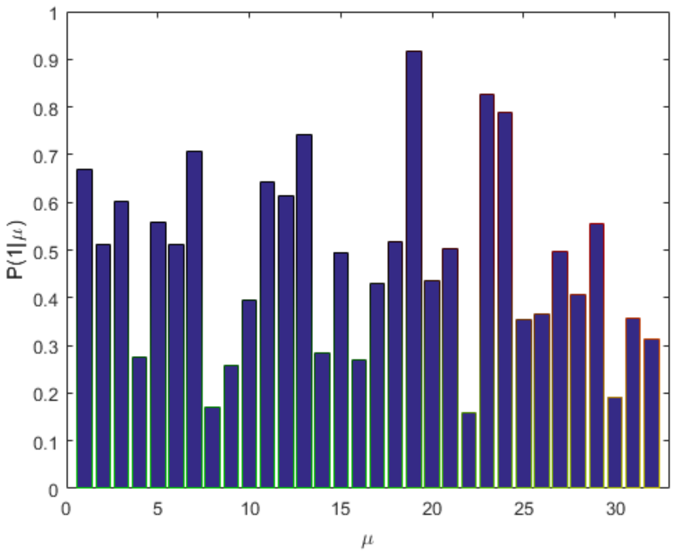

Figure 3: Histogram of conditional probability of choosing the bid 1 given a specific history for $\alpha = \frac{2^5}{51} > \alpha_c$.

3 Grand Canonical Minority Game

The Minority Game obviously has some major deficiencies as a market model. For starters it assumes the agents will always participate on the market regardless of how well their strategies are performing. Also, it does not show a fat tailed price return but has Gaussian fluctuations. For this purpose The Minority Game has been expanded upon in different ways, one of them being the so called Grand Canonical Minority Game which we will study in some detail in this chapter [8].

3.1 Description of the game

The Grand Canonical Minority Game is a development of the basic Minority Game with the aim of making it more like a real market. In this model the agents can opt not to participate in the game if their strategies have not been performing well enough. This additional option for the agents gives rise to fluctuating volumes on the market. It is also how the version of the game has gotten its name: in the grand canonical ensemble in statistical mechanics the number of particles varies in the observed system.

In our work the option not to participate is thought of as the second strategy which means the agents are only equipped with one regular strategy. To build on this further we introduce an external cost, ϵ, in the score updates for the regular strategy. This cost can be seen as what would have happened with the agent's money if the agent opted to put his money in a bank account with steady interest rate. This is equivalent with giving the agent an inactive strategy which get to have the score ϵt after t time steps. In addition we add producers to the model who must participate all the time and they are only equipped with one strategy each. This version of the game preserves some of the major feature from the basic game, most importantly the two phased nature with an unpredictable respective a predictable phase [9].

In the Grand Canonical Minority Game we use linear payoff in the update scheme. The score update $U(t+1)$ for an agents strategy for the Grand Canonical Minority Game will thus be

$$U_i(t+1) = U_i(t) - \epsilon - a_i^{\mu(t)} A(t), \tag{3}$$

for some $\epsilon > 0$ and $A(t)$ as

$$A(t) = \Omega^{\mu(t)} + \sum_{i=1}^{N_s} a_i^{\mu(t)} \phi_i(t), \tag{4}$$

where Ω^μ is the net contribution from the producers for the history μ and $\phi_i(t)$ is 0 or 1 depending on whether the agents participates or not at time t. In this version of the game we call the number of agents N_s and will sometimes refer to them as speculators. To determine whether or not an agent should participate we have the following rule

$$\phi_i(t) = \begin{cases} 1 & : U_i(t) \geq 0 \\ 0 & : U_i(t) < 0 \end{cases}.$$

Again we update the scores for every agent, including those that are inactive. This model clearly poses a problem in the way we defined the history in the previous game seeing as there will not necessarily be a winning bid in every time step. This causes issues in the simulation since the agents should base their decisions on the m last time steps winning bid. However, it turns out that the game qualitatively remains intact with the use of a randomized history in each time step [10][11]. The main features of the game remains intact. Most crucial is that all agents react to the same event, rather than that event being what actually happened. We will use this in our work. Another way of solving this problem would be to simply toss a coin in the event of a non-winning bid.

3.2 Producers

The reason producers are added to the model is because otherwise we would find ourselves in a situation without any active agents after a while [12]. The producers have strategies, much like the agents themselves, that they follow. If we have N_p producers we have

$$\Omega^\mu = \sum_{i=1}^{N_p} a_i^\mu \tag{5}$$

as their total contribution. Since Ω^μ is a sum of N_p uniformly distributed variables with values of either 1 or -1 its mean is obviously $\langle \Omega^\mu \rangle = 0$. The variance is, just as for the attendance in the Minority Game, $\langle (\Omega^\mu)^2 \rangle = N_p$. Thus, for large N_p by the Central Limit Theorem, $\Omega^\mu \sim \mathcal{N}(0,N_p)$. The vector $\vec{\Omega} = (\Omega^1,\Omega^2,...,\Omega^P)$ for $\mu = 1,2,...,P$ is a vector with P elements all being in the range [-N_p,N_p], where only every second value is attainable. We see this from Equation 5, where if one agent that bid 1 changes to -1 the change will be two steps and every value cannot be attained. The event μ that causes the largest absolute net contribution from the producers will be referred to as extreme event and the corresponding time step is called critical.

3.3 Behaviour for different m

One interesting feature of this game is the behaviour we get when simulations are made with a fairly small strategy space, especially if one manually create the vector of the producers. What we do when manually creating the vector is simply making sure that there is a rather large extreme event. We see a typical run of the game done in that way in Figure 4, where we see the evolution of strategy scores in time. First we note the chaotic first phase where there are very large fluctuations. After a while, however, we note that the agents divide into two distinctly different groups: one that is around zero and one that goes down continuously. This is an important feature of the game in this regime: the agents that are anti correlated against the producers extreme event (which is the history that has the combined biggest contribution of the producers) will harbour around zero points whereas the agents correlated to that event will continue downward. Correlation to an event is making the same bid as the sign of the combined contribution of the producers.

Another interesting feature is the peaks, or 'booms' as they would be called on an actual market, that appear in the simulation. We can see they happen fairly regularly and seem to fizzle out fast when they do. This is an effect of the small strategy space: we will have a fair few agents answering the same to quite a few histories. This, again, is crowd effects, much like those we saw in previous section. The time step where the peak fizzles out is a critical time step where all the agents participating in the game are anti correlated to the producers combined contribution, and they will tip the attendance over to the other sign making them all to end up with the non-winning bid. These effects, however, seems to only appear in simulations of the model with small strategy space which is reasonable considering the correlation between agents bids will lessen as P grows.

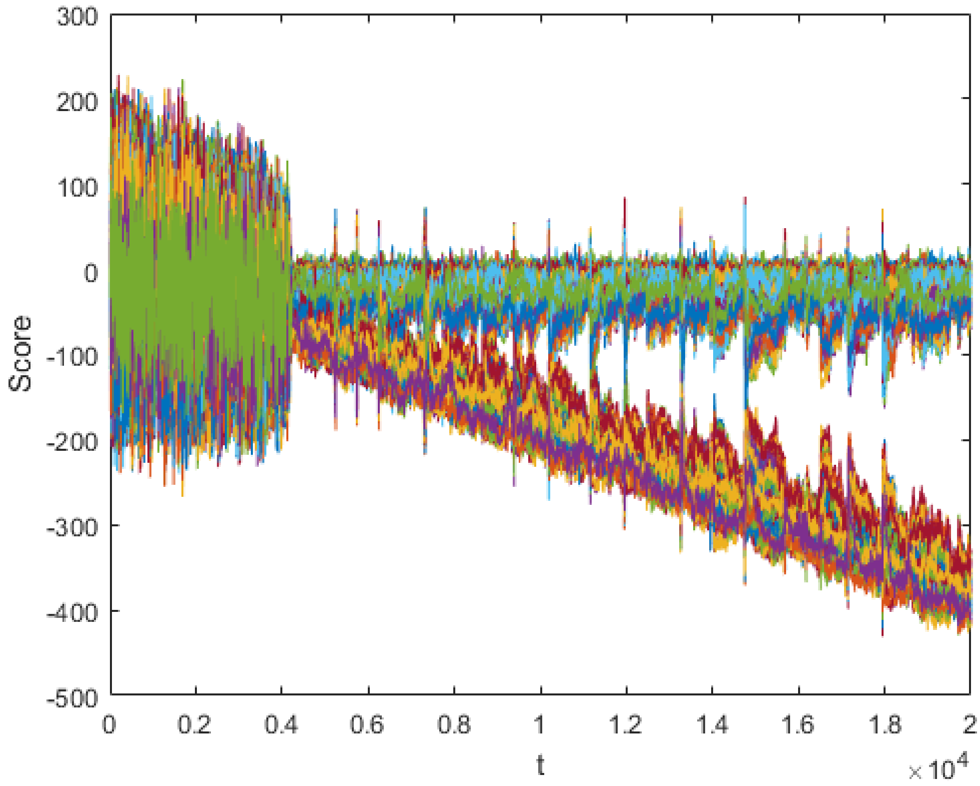

Figure 4: Here we can see a realization of the game for $m = 4$ and $N_s = 20P$. We can clearly see the strategies dividing in two distinct groups. We also note spikes in strategy scores which happens fairly regularly.

For larger values of m these effects disappear. A typical run of the game can instead look like in Figure 5. The initial phase is not as obvious here as it was for smaller m. The length of this phase is governed by how many agents we have compared to P as well as how large ϵ is. We see that the strategies divide and as expected more tends to go downward than upward. This choice of parameters might not seem as interesting on the surface as the one above but nevertheless our analysis of the game will be in this regime. One of the reasons for this is that it will reflect a market where agents acts more independently from each other. Looking at Figure 5 again we see that the agents might be classified into different groups as discussed further in the next section.

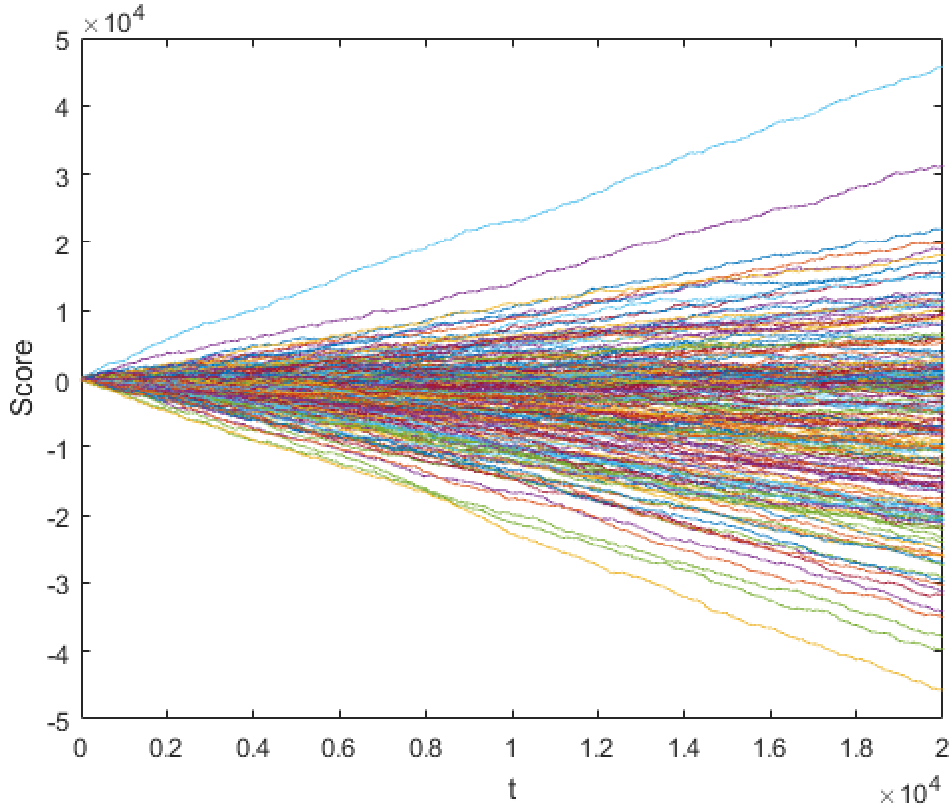

Figure 5: Here we can see a realization of the game for $m = 8$ and $N_s = P$. We see that we do not have the same effects as above. We note that the strategies generally seem to loose points.

3.4 Different types of agents

As touched upon, in simulations of this model we have detected that different kind of agents appear. In essence we have three types of agents that can be put into three distinct groups

- Negatively frozen agents. These agents will never take part in the game, which simply means their strategy never is prosperous enough.

- Positively frozen agents. These will be the agents that takes part in the game at all time steps.

- Fickle agents. The fickle agents are those that takes part in the game on occasion. This means that their strategy score will be around zero, having tops over and dips under.

We can see typical behaviour of such agents in Figure 6. We note that after a short initial phase they behave as we expect them to. To clarify a bit how we differentiate the different types of agents in numerical simulations we say that after a long time t (for example $t \approx 10^6$) we expect initial disturbances to have vanished and thus the agents to behave as described above. If we define m_i as the fraction of time steps agent i participates in the game, we expect that $m_i \to 0$ for negatively frozen agents and $m_i \to 1$ for positively frozen agents as $t \to \infty$.

Since we will have fickle agents in the game the amount of active agents will fluctuate. This means that an interesting quantity to study is the amount of active agents, which is defined as $N_{act}(t) = \sum_{i=1}^{N_s} \phi_i(t)$. Much like the basic Minority Game will the average attendance of this game be 0. The variance, however, is again more interesting. If we calculate this a little naively, the variance takes the form $\sigma^2(t) = N_p + N_{act}(t)$. This is what it would look like if the bids of the speculators were uncorrelated to those of the producers. However in the development of our model we will see that the bids of the active agents are in fact, on average, anti correlated to the combined contribution of the producers leading to a smaller variance.

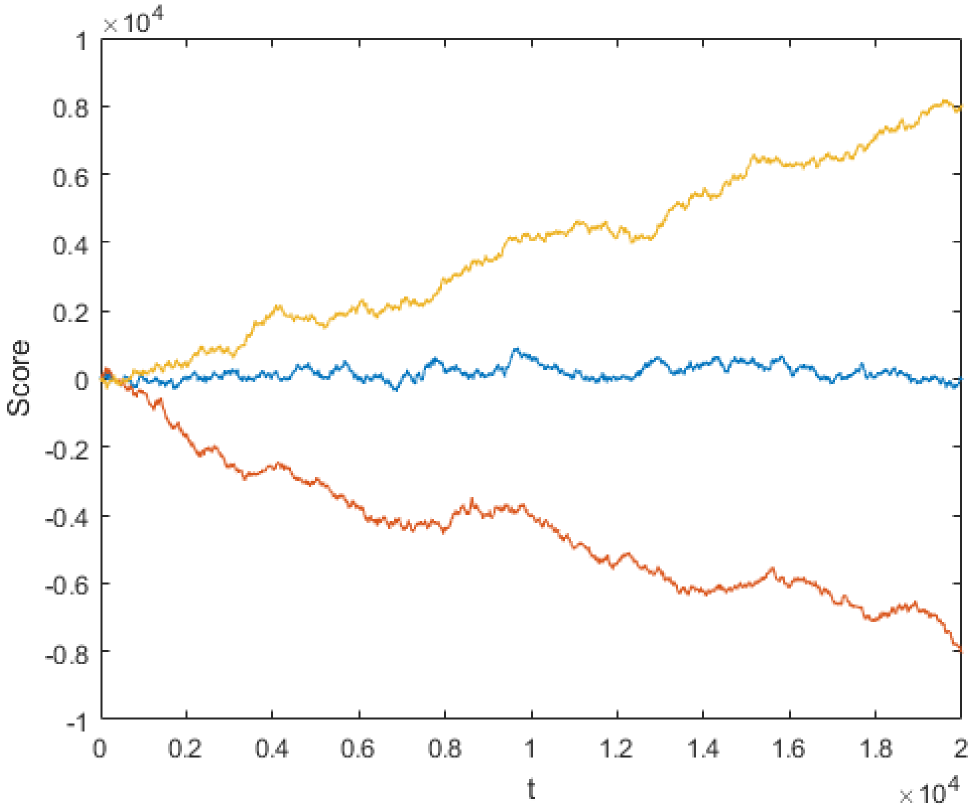

Figure 6: Here we have taken out the three agents with different characteristics from the above simulation. We see that after a time, the yellow agent is always active (score over zero), the blue one are occasionally active (score around zero) and the red one is never active (score under zero).

3.5 Reduced parameters

Analyzing these kind of problems in statistical mechanics, one are often interested in finding the solution in the thermodynamic limit. The thermodynamic limit, in this problem, is when we let $P \to \infty$ whereas we keep the reduced parameters $n_s = \frac{N_s}{P}$ and $n_p = \frac{N_p}{P}$ constant. This means that for our statistical analysis of the game we will be interested in the parameters n_s and n_p instead of the actual amount of agents and producers. Considering we do not have infinite computer resources, we will study the game for moderately large values of m, but will not exceed $m = 10$. However, we remember that it is equivalent to $2^{2^{10}}$ distinct strategies, which means we have a very large strategy space.

3.6 Stylized facts of the model

One of the reasons why this model is interesting to study is that it reproduces some 'stylized facts' from the real market. Stylized facts are something that has been observed empirically so many times that it is accepted as truths. The stylized facts that the model reproduces include a fat tailed price return as well as volatility clustering [9]. Volatility clustering, for example, can easily be seen in figure 4 around the 'booms'. Generally we get these kinds of fat tails for smaller systems, as the analysis of Challet and Marsili shows [13]. The study of such models is a tool to understand how macroscopic effects are caused by the microscopic dynamics of each agent [3]. However the regime we will be looking at will not show these fat tailed price return but rather a normal return.

3.7 Analytical Model

The fact that we, as we mentioned earlier, use randomized histories turns out to simplify an analytical approach. Otherwise the explicit time feedback induced by the memory would have made the system non-Markovian and thus an analytical approach would have been much harder. However with randomized histories we bypass this problem as the agents now react to noise in each time step and we arrive at a stochastic and Markovian problem [14].

In our analytical study of the game we are primarily interested in studying and finding the agents mean step length $\frac{1}{P}\sum_{\mu=1}^{P}\langle\Delta_i(t)\rangle$, which is to say step length averaged both in time and over the different histories. Here we define the step length as $\Delta_i(t) = U_i(t) - U_i(t-1)$. We will denote time averages as $\langle\cdots\rangle$ and in the continuum limit history averages as $\overline{(\cdots)}$. We find a distribution for mean step sizes from which we can add the negative bias induced by active agents using their strategy. The negative bias arise from Equation 3 and 4, where active agents will self-interact with themselves. Positively frozen agents will be those that are able to on average overcome the bias of using their strategy in each time step whereas fickle are those that when active can not overcome their own bias. Proceeding, with these individual averaged step sizes we will be able to find the score distributions of the different types of agents by solving the master equation on an integer chain. From the master equation we get $P_i(x, t)$ - the probability to find agent i with score x at the time t. Since simulations are done over many agents it is of interest to find the full score distribution $P(x, t) = \int P(\Delta_i)P_i(x, t)d\Delta_i$. Further, the model also makes it possible to quantify the anti correlation of the active agents choices of bids towards the producers. The main feature, however, is that we are able to calculate the fractions of the different types of agents in the game for different values of n_s and n_p. The model will then be compared to numerical simulations where we numerically calculate fractions as well as the full score distribution.

4 Statistical Model

In this section we will derive a statistical model for the Grand Canonical Minority Game in the regime $n_s \leq 1$. With the model we will be able to get fractions of different types of agents as well as performing a random walk on an integer chain to find the full score distribution.

4.1 Distribution of mean step sizes

In this section we derive the distribution for the mean step size of the agents as well as expressions for fractions of different types of agents. Some of the analytic calculations will be left out of this part and will be presented in Appendix A in more detail.

We start by repeating some basics from earlier sections. From the update scheme in Equation 3, we define the step length as $\Delta_i(t) = U_i(t) - U_i(t-1)$. This means the step length Δ for a particular agent i at the time t is written as

$$\Delta_i(t) = -a_i^\mu(\Omega^\mu + \sum_{i=1}^{N_s} \phi_i(t) a_i^\mu) - \epsilon.$$

We want to proceed to find the time and history averaged step sizes for each agent. The first step in doing this is by removing the time-dependence of this quantity. Hence we start by taking the time average of $\Delta_i(t)$ which means we get an expression that only depends on the history μ and arrive at

$$\Delta_i^\mu = -a_i^\mu(\Omega^\mu + \sum_{i=1}^{N_s} m_i a_i^\mu) - \epsilon,$$

where $m_i = \frac{1}{T}\sum\limits_{t=1}^{T} \phi_i(t)$ is the the average amount of times an agent participates in the game. Obviously $0 \leq m_i \leq 1$. As discussed previously we have different type of active agents: the positively frozen agents which will be modeled as to have $m_i = 1$ and fickle agents where $0 < m_i < 1$. In order to arrive at expressions for the mean step length we need to make a number of assumptions. We start by making this simple assumption: the fickle and frozen agents will be modeled to be drawn from different distributions. This means we can write the step length by separating between these agents as

$$\Delta_i^\mu = -a_i^\mu(\Omega^\mu + \sum_{i=1}^{f^+ N_s} a_i^\mu + \sum_{j=1}^{k f^0 N_s} a_j^\mu) - \epsilon,$$

where f^+ the fraction positively frozen agents, f^0 is the fraction fickle agents and k is the fraction active fickle agents on average. What we have done here is simply assume that the choices a_i^μ will be drawn from different distributions for fickle respective frozen agents. We also realize that since all active fickle agents are drawn from the same distribution we can sum their m_i together into kf^0N_s.

From now on we call the sum of the contributions for the positively frozen agents $X^\mu = \sum_{i=1}^{f^+N_s} a_i^\mu$, and the corresponding sum for fickle agents $Y^\mu = \sum_{j=1}^{kf^0N_s} a_j^\mu$. We are interested to find the distribution of these two objects, and more specifically their dependence on Ω^μ. To do this we proceed by making an additional assumption: if we look at the joint distribution $a_i^\mu\Omega^\mu$, we assume that the positively frozen agents, over all histories, will be the most anti correlated to the net contribution of the producers. In the same way we assume the negatively frozen agents to be the most correlated. These objects are also assumed to be independent.

We will also drop the index i on all these objects and instead consider these variables to be drawn from different distributions. To do this we need to represent these objects with probability distributions. If we look at a^μ, this is a quantity that will only take the value 1 or -1 for individual agents. In the proceeding analysis we will instead represent a^μ by assuming it is normal with mean 0 and variance 1 (since a_i^μ has mean 0 and variance 1). This implies that the joint object $a^\mu\Omega^\mu$ has mean 0 and variance N_p. Furthermore, if we look at the object

$$\sum_{\mu=1}^{P} a^\mu\Omega^\mu = \vec{a}\cdot\vec{\Omega},$$

it will be the sum of P such independent random variables. This implies, by the Central Limit Theorem, for large P, that $\vec{a}\cdot\vec{\Omega}$ is normally distributed with mean 0 and variance N_pP. From this distribution we will now be able to find the positively frozen agents according to our previous assumption. Our reasoning is this, since we expect these agents to be the most anti correlated to the producers they will be on the left hand side of the distribution and there will exist a c such that $\vec{a}\cdot\vec{\Omega} < -c$ for all positively frozen $\vec{a}$. This means can calculate f^+ to be

$$f^+ = \int_{-\infty}^{-c} \mathcal{N}_x(0,\, N_pP)\,\mathrm{d}x = \frac{1}{2}\operatorname{erfc}\left(\frac{c}{\sqrt{2N_pP}}\right), \tag{6}$$

where $x = \vec{a} \cdot \vec{\Omega}$ and $\mathcal{N}_x(\mu,\, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$. In the same manner we can calculate the fraction of negatively frozen agents for some other constant d to be

$$f^- = \int_d^\infty \mathcal{N}_x(0,\, N_p P)\, \mathrm{d}x = \frac{1}{2} + \frac{1}{2}\,\mathrm{erf}\left(\frac{d}{\sqrt{2N_p P}}\right).$$

The fraction fickle agents are then calculated from these two to be $f^0 = 1 - f^+ - f^-$. Further, as previously mentioned, we are interested to find a measure of the anti correlation of a^μ in relation to Ω^μ for the active agents. Since we assumed the positively frozen agents to be anti correlated with Ω^μ we will have

$$\mathrm{P}_{a^\mu|\Omega^\mu} = \mathcal{N}_{a^\mu}(\mu_1,\, 1),$$

for some μ_1 which we assume is linearly dependent on Ω^μ, that is $\mu_1 = -g_+\Omega^\mu$. To find this g_+ we proceed to take the conditional average of the positively frozen agents over $\vec{a} \cdot \vec{\Omega}$:

$$\overline{(\vec{a} \cdot \vec{\Omega})}_{f^+} = \frac{1}{f^+} \int_{-\infty}^{-c} x\mathcal{N}_x(0,\, N_p P)\, \mathrm{d}x = -\frac{1}{f^+}\sqrt{\frac{N_p P}{2\pi}} e^{-\frac{-c^2}{2N_p P}}.$$

We then calculate the same conditional mean in a different way:

$$\overline{(\vec{a} \cdot \vec{\Omega})}_{f^+} = P\overline{(a^\mu \Omega^\mu)}_{f^+} = P \int a^\mu \Omega^\mu \mathrm{P}(a|\Omega)\mathrm{P}(\Omega)\, \mathrm{d}a\, \mathrm{d}\Omega = -P g_+ N_p.$$

With these two equations we can solve for g_+. In the same manner we assume that the mean for fickle agents are linearly dependent on Ω^μ and we can solve for fickle agents to attain their anti correlation constant g_0 to the producers. We thus arrive at

$$\mathrm{P}(X^\mu|\Omega^\mu) = \mathcal{N}(\text{-}g_+\Omega^\mu f^+ N_s,\, f^+ N_s), \qquad \mathrm{P}(Y^\mu|\Omega^\mu) = \mathcal{N}(\text{-}g_0\Omega^\mu k f^0 N_s,\, k f^0 N_s),$$

with $g_+ = \frac{1}{f^+}\sqrt{\frac{1}{2\pi N_p P}} e^{-\frac{-c^2}{2N_p P}}$ and $g_0 = \frac{1}{f^0}\sqrt{\frac{1}{2\pi N_p P}} (e^{-\frac{-c^2}{2N_p P}} - e^{-\frac{-d^2}{2N_p P}})$.

Moving forward we want to find the history averaged step length, $\frac{1}{P}\sum_\mu \Delta^\mu$. To do this we repeat that we will have a uniform history distribution because we are randomizing the histories. In our expression for Δ^μ we also insert a bias term explicitly according to

$$\Delta^\mu = -a^\mu(\Omega^\mu + X^\mu + Y^\mu) - \theta(x) - \epsilon, \quad \theta(x) = \begin{cases} 1 & : x \geq 0 \\ 0 & : x < 0 \end{cases}.$$

Since the bias term, $\theta(x)$, is history independent we have that $\mu_{bias} = \frac{1}{P}\sum_{\mu=1}^{P}\theta(x) = \theta(x)$ and its variance will be $\sigma^2_{bias} = \frac{1}{P}\sum_{\mu=1}^{P}\theta(x)^2 - \theta(x)^2 = 0$. We thus turn our attention to the first part of this expression which we call δ^μ ($\delta^\mu = -a^\mu(\Omega^\mu + X^\mu + Y^\mu)$). This object will determine the agents mean step size distribution. For large P we assume independence between distributions and we thus easily attain $\frac{1}{P}\sum_{\mu=1}^{P}\delta^\mu = 0$ (since a^μ is a zero-mean variable). The variance for δ^μ is calculated through

$$\sigma^2_\delta = \frac{1}{P}\int (a^\mu(\Omega^\mu + X^\mu + Y^\mu))^2 \mathrm{P}(\Omega)\mathrm{P}(a)\mathrm{P}(X|\Omega)\mathrm{P}(Y|\Omega)\,\mathrm{d}X\,\mathrm{d}Y\,\mathrm{d}\Omega\,\mathrm{d}a\,, \tag{7}$$

in the continuum limit. This calculation is left for Appendix A but we get

$$\sigma^2_\delta = n_p(1 - N_s(f^+g_+ + kf^0g_0))^2 + f^+n_s + kf^0n_s,$$

and has thus derived the distribution for the individual averaged mean step sizes for the agents. We can now calculate the fraction positively frozen agents from the distribution over mean step sizes through

$$f^+ = \int_{1+\epsilon}^{\infty} \mathrm{P}(x)\,\mathrm{d}x\,, \tag{8}$$

where $\mathrm{P}(x) = \mathcal{N}(0,\, \sigma^2_\delta)$. This simply means we say that the positively frozen agents will have a mean step size over $1+\epsilon$ and will thus be able to on average overcome their own induced bias as well as the external cost. In the same way the negatively frozen agents will be from the same distribution but with $x < \epsilon$ and the fraction fickle in between.

$$f^- = \int_{-\infty}^{\epsilon} \mathrm{P}(x)\,\mathrm{d}x \quad , \quad f^0 = \int_{\epsilon}^{1+\epsilon} \mathrm{P}(x)\,\mathrm{d}x$$

All these equations are solved self-consistently in the model by setting that the fraction of a certain type of agent must be equal to the same fraction calculated in another way, for example Equation 6 equals Equation 8.

In addition we can find the conditional expectations of the mean step lengths of different types of agents from this distribution and with these we will be able to find the full score distributions of all agents.

4.2 Score distributions

As we now have characterized each agents mean step size we can now ask ourselves what probability there is to find a certain type of agent with score x at time t. To answer this question we need to find the full score distributions. This, in turn, is done by first solving the so-called master equation on an integer chain.

To start we need to find the jump probabilities on the chain. The jump probabilities are the probabilities to jump from one score x to another score x' in the next time step. For this purpose we repeat that we will have different mean step sizes for different agents and we note that $A(t)$, which will determine the hopping length, has the variance $\sigma^2 = P\sigma_\delta^2$. To see this, Equation 7 gives us the variance, σ_δ^2, of the mean step size and the variance of the step size would be that expression multiplied with P. The variance of the step size is equal to the variance of the attendance considering $\mathrm{P}(a^\mu) = \mathcal{N}(0,\, 1)$ and the assumption of independence of distributions. Further, since the components in the attendance are all normal-distributed and the hopping is decided by the attendance, the jump probabilities should be well approximated by a normal distribution

$$\mathrm{p}_{x \to x'} = \mathcal{N}_{(x'-x)}(\Delta_{mean},\, \sigma^2),$$

where Δ_{mean} is the averaged step length for the different types of agent given by

$$\underbrace{\Delta_{f^+} = \frac{1}{f^+}\int_{1+\epsilon}^{\infty} \mathrm{P}(x)x\,\mathrm{d}x - 1 - \epsilon}_{\text{average step length positively frozen agents}} \quad , \quad \underbrace{\Delta_{f^-} = \frac{1}{f^-}\int_{-\infty}^{\epsilon} \mathrm{P}(x)x\,\mathrm{d}x - \epsilon}_{\text{average step length negatively frozen agents}}$$

$$\underbrace{\Delta_{+} = \frac{1}{f^0}\int_{\epsilon}^{1+\epsilon} \mathrm{P}(x)x\,\mathrm{d}x - 1 - \epsilon}_{\text{average step length fickle agents } x\geq 0} \quad , \quad \underbrace{\Delta_{-} = \frac{1}{f^0}\int_{\epsilon}^{1+\epsilon} \mathrm{P}(x)x\,\mathrm{d}x - \epsilon}_{\text{average step length fickle agents } x<0}$$

with $\mathrm{P}(x) = \mathcal{N}(0,\, \sigma_\delta^2)$. The master equation then takes the form

$$\mathrm{P}_x(t+1) = \sum_{x'} \mathrm{p}_{x' \to x}\mathrm{P}_{x'}(t),$$

where $\mathrm{P}_x(t)$ is the score distribution. The master equation gives the probability to find an agent with score x at time t by adding together the possibilities to jump to that score from every other score. This can be solved in the continuum limit for fickle agents, and we find the score distribution. For fickle agents we assume their score distribution to be stationary, such that $\mathrm{P}_x(t) = \mathrm{P}_x$. This seems like a reasonable assumption since fickle agents needs to have around zero score and thus cannot have a time-dependent drift from zero. This leads to

$$\mathrm{P}(x) \sim e^{\mp 2|\Delta_\pm| x/\sigma^2}, \tag{9}$$

where $\Delta_\pm$ is the conditional averaged step size for active respective inactive fickle agents. Details of this derivation can be seen in Appendix A.

For frozen agents we realize that they will, on average, drift further from 0 for each time step. The same problem is solved by Granath and Perez-Diaz in [15] as well, where they used the derivation of Marsili and Challet [16]. The expression is derived by solving the Fokker-Planck equation. It will be a diffusion with a drift and takes the form

$$\mathrm{P}(x,\, t) = \mathcal{N}_x(\Delta_{f^+, f^-} t,\, \sigma^2 t),$$

where we have different conditional step sizes for positively respective negatively frozen agents.

From these different probability distributions we can find the full score distributions for the agents by integrating them over the distribution of respective step sizes. The full distribution will be $\mathrm{P}(x,\, t) = \int \mathrm{P}(\Delta_i)\mathrm{P}_i(x,t)\,\mathrm{d}\Delta_i = \mathrm{P}_{f^0}(x) + \mathrm{P}_{f^+}(x,t) + \mathrm{P}_{f^-}(x,t)$. The component for fickle agents is computed as

$$\mathrm{P}_{f^0}(x) = \int_{\epsilon}^{1+\epsilon} \frac{\mathcal{N}_z(0,\, \sigma_\delta^2) 2(\sigma^2)^{-1} e^{2(z-\epsilon-\theta(x))x/\sigma^2}}{(z-\epsilon)^{-1} - (z-1-\epsilon)^{-1}}\, \mathrm{d}z\,,$$

where we have normalized the expression in Equation 9 over x. For frozen agents we have

$$\mathrm{P}_{f^+}(x,\, t) = \int_{1+\epsilon}^{\infty} \mathcal{N}_z(0,\, \sigma_\delta^2)\mathcal{N}_x(t(z-1-\epsilon),\, \sigma^2 t)\,\mathrm{d}z$$

$$\mathrm{P}_{f^-}(x,\, t) = \int_{-\infty}^{\epsilon} \mathcal{N}_z(0,\, \sigma_\delta^2)\mathcal{N}_x(t(z-\epsilon),\, \sigma^2 t)\,\mathrm{d}z$$

Additionally, from the expression for $\mathrm{P}_{f^0}(x)$ we find the active fraction of fickle agents k to be

$$k = \frac{\int_0^\infty \mathrm{P}_{f^0}(x)\,\mathrm{d}x}{\int_{-\infty}^{\infty} \mathrm{P}_{f^0}(x)\,\mathrm{d}x}.$$

Finally we, as Granath and Perez-Diaz did in [15], can integrate and average the full distribution $\mathrm{P}(x,\, t)$ over a time window to attain the probability to find an agent with a certain point x between time t_0 and t_1.

5 Results

We now turn to comparing our model to numerical simulation to see how well it works. As previously mentioned, the most important features of the model includes calculating the fraction of the different kinds of agents as well as finding the full score distribution. We will try our model for different values of the parameters n_s, n_p and m.

5.1 Fractions of agents

In Figures 7, 8, 9 and 10 we see simulations and corresponding model values for the different fractions with different values of n_s, n_p and m. All numerical data points are averaged over 20 runs and the bars are the standard deviations of the simulations. We note that our model recreates the basic features of how the different fractions behave for for different values of n_p. Thus the model is in pretty good qualitative agreement with simulations. In particular the model seems to fit very well for $n_s = 0.5$. We also note that the deviations in simulations seems to lessen when m gets larger. This is very well in line with what one would have thought as larger m means that we have larger strategy space and thus the randomization of strategies in the initiation will have lesser impact. We note that for $n_s \leq 1$ we have fairly good quantitative agreement as well. However, as we can see in in Figure 11 the model differ quite considerably quantitatively as well as qualitatively from numerical simulations for $n_s = 2$. For larger n_s these deviations are even bigger and we conclude that the model does not work at all in these regimes.

In Figure 12 we see how correlated the different types of agents are towards the producers. We see that our assumption turns out to be in good agreement with numerical simulations.

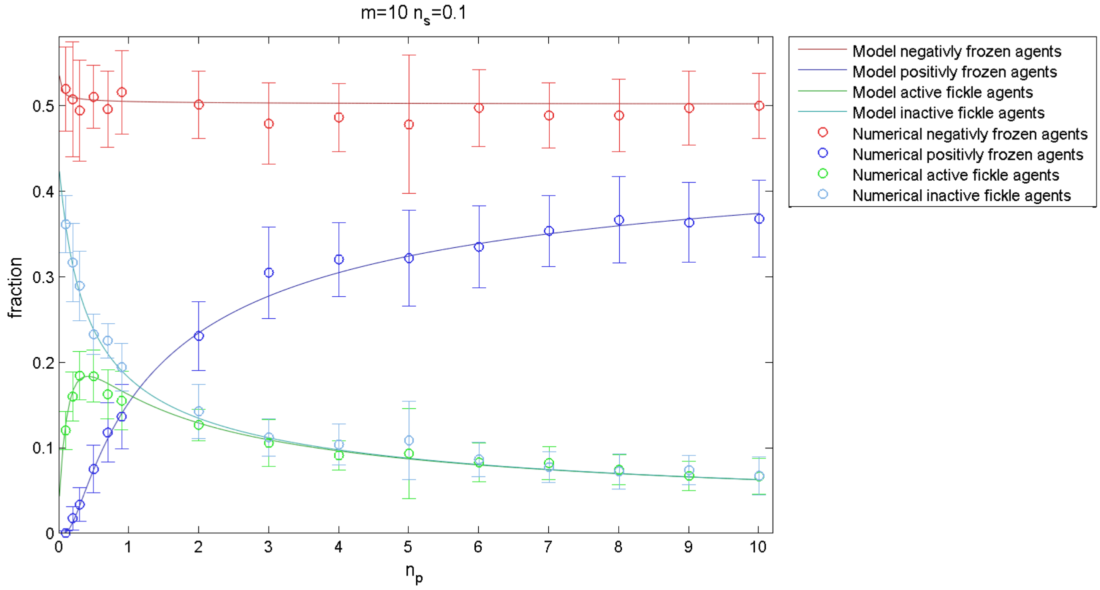

Figure 7: With parameters $m = 10$ and $n_s = 0.1$, we can see how well the characteristics of the model and numerical simulations correspond. Numerical data is the average over twenty simulations and the standard deviation is plotted to visualize the deviations.

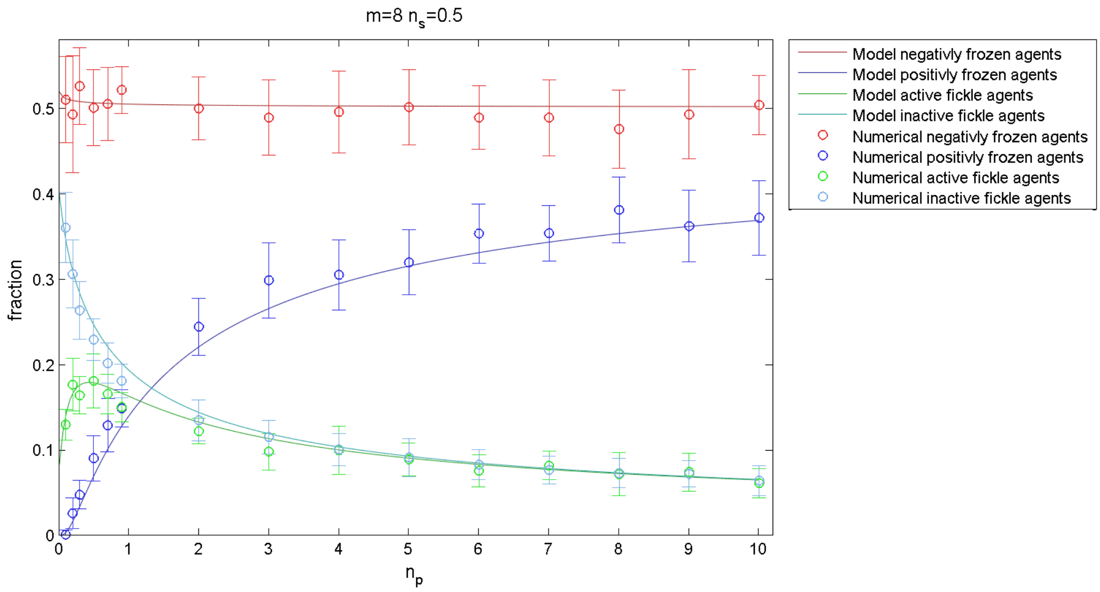

Figure 8: With parameters $m = 8$ and $n_s = 0.5$, we can see how well the characteristics of the model and numerical simulations correspond. Numerical data is the average over twenty simulations and the standard deviation is plotted to visualize the deviations.

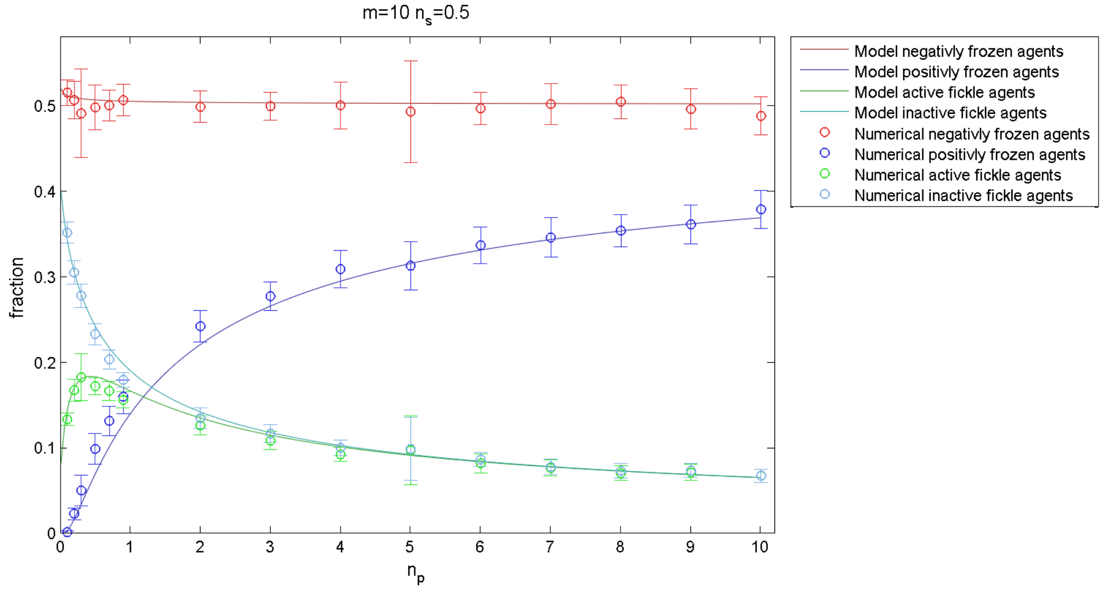

Figure 9: With parameters $m = 10$ and $n_s = 0.5$, we can see how well the characteristics of the model and numerical simulations correspond. Numerical data is the average over twenty simulations together with standard deviation. With larger m the simulations get smaller standard deviations.

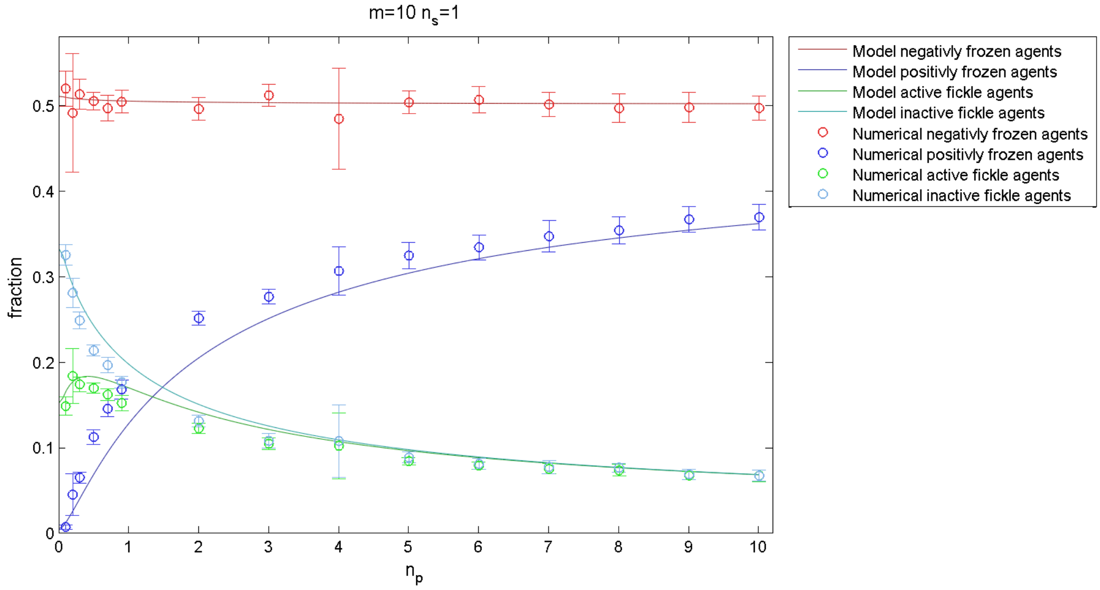

Figure 10: With parameters $m = 10$ and $n_s = 1$, we can see how well the characteristics of the model and numerical simulations correspond. Numerical data is the average over twenty simulations together with standard deviation.

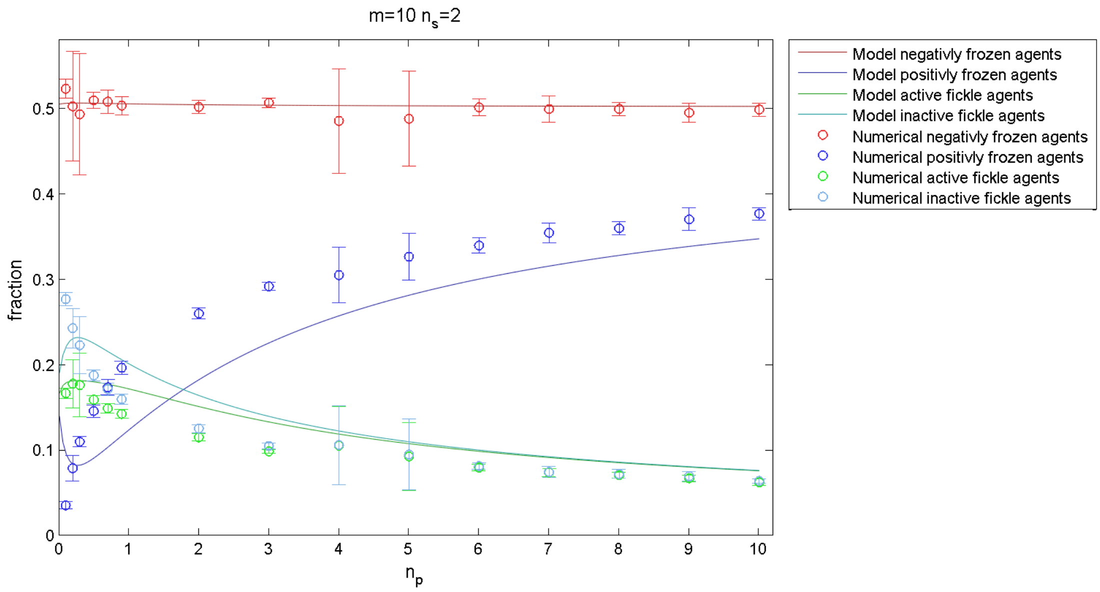

Figure 11: With parameters $m = 10$ and $n_s = 2$, we can see how well the characteristics of the model and numerical simulations correspond. Numerical data is the average twenty simulations together with standard deviation. We note that for $n_s = 2$ the model is quite a bit off, and you can clearly see how it is qualitatively different for small n_p.

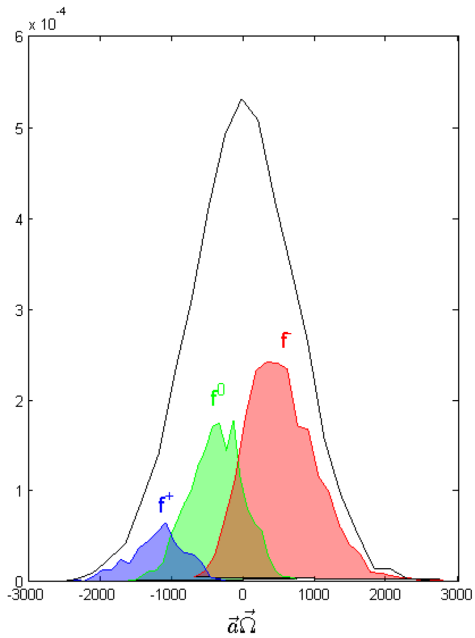

(a) Numerical fractions of agents.

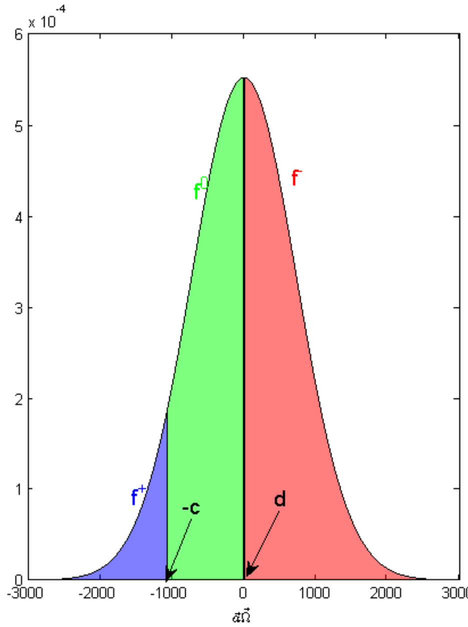

(b) Modelled fractions of agents, where c and d are the limits of integration.

Figure 12: Distribution $\mathrm{P}(\vec{a}\vec{\Omega})$ from simulation as well as theoretical, with parameters $m = 10$ and $n_s = n_p = 0.5$. We note that our assumption of correlation seems reasonable and we see that the conditional averages roughly will be the same in both cases.

5.2 Score distribution for fickle agents

In Figures 13a and 13b we see the score distributions for fickle agents from a simulation where we have looked at the distribution at two different times. We see that these distributions looks fairly similar and we conclude that our assumption that it is time-independent is very reasonable.

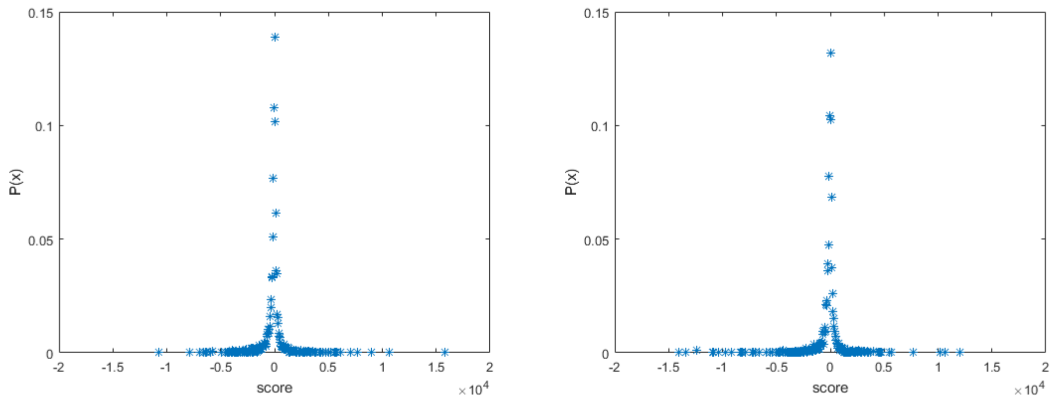

(a) Numerical score distribution for fickle agents at time $t = 3.5 \cdot 10^6$.

(b) Numerical score distribution for fickle agents at time $t = 4 \cdot 10^6$.

Figure 13: Numerical score distribution for fickle agents at different times t.

5.3 Full score distribution for all agents

In Figure 14 wee see the full score distribution, integrated over a time window, of all agents compared to numerical data for $m = 8$, $n_s = 0.5$ and $n_p = 2$ ($\frac{1}{t_1 - t_0} \int_{t_0}^{t_1} \mathrm{P}(x,t)dt$, where $\mathrm{P}(x,t)$ is the full distribution as seen above). We note that the theoretical distribution overall agrees very well with simulations, however as we see in Figure 15 it seems to under predict the probability for agents to have scores very close to zero. Note that the distribution is a bit skewed to the left - this is because we have fewer positively frozen agents than negatively frozen.

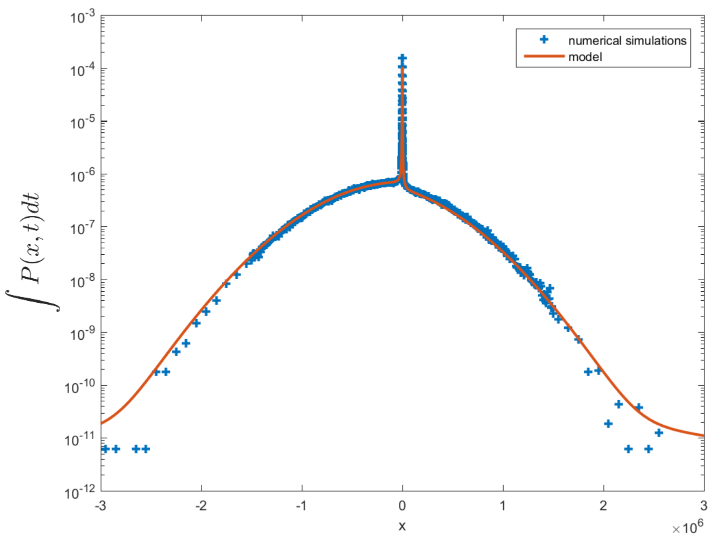

Figure 14: Full score distribution integrated over a time window $t_0 = 4 \cdot 10^5$ to $t_1 = 5 \cdot 10^5$, $\frac{1}{t_1 - t_0} \int_{t_0}^{t_1} \mathrm{P}(x,t)dt$, with parameters $m = 8$,$n_s = 0.5$ and $n_p = 2$. The agreement to model values is very good. To note, there is only a few values for the tails, and the model under predicts close to $x = 0$. The numerical data is from 12600 runs with the above parameters.

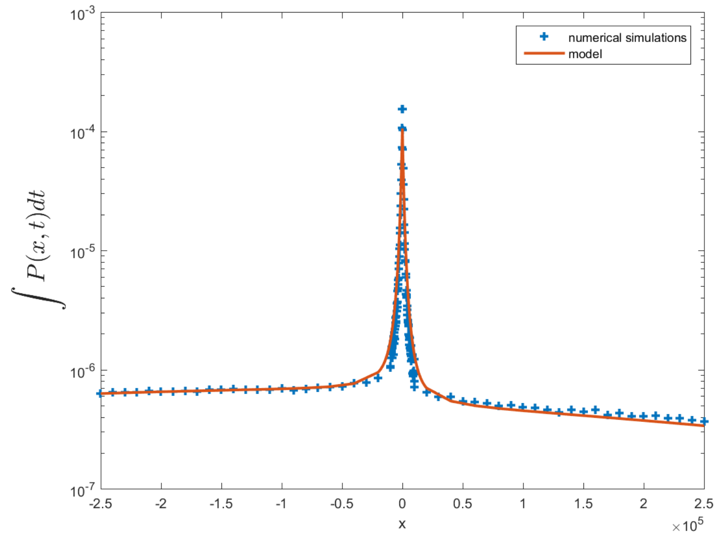

Figure 15: The full score distribution zoomed in around $x = 0$. We clearly see that it under predicts the probability of very small values of x.

6 Summary and Discussion

In the first chapter we started the work by familiarizing ourselves with Minority Game and recreating some of its basic features. These features included showing how the volatility σ^2/N depends on $\alpha = 2^m/N$ and how there is a phase transition between a unpredictable and predictable phase. In the next chapter we moved on to study the Grand Canonical Minority Game where we first went through the basics of the game and why it is an interesting model. We introduced three different types of agents and how to decide in which group to place each agent in numerical simulations. Further, reasons are given for the choice of regime we study. We then turn to describing what has been studied in our analytical model of the game.

In the following chapter we go through our derived statistical model in detail, describing how we arrived at expressions for the mean step size for agents as well as expressions for fractions of the different types of agents. Having done that we turned our attention to deriving full score distribution for the agents by solving the master equation on an integer chain and then integrating those expressions over the mean step size distribution. In the last chapter we looked at how our model compared with numerical simulations.

Now we turn to discuss how well our model works as well as possible improvements for future works. We have derived an analytical model that both qualitatively and quantitatively are in fairly good agreement with numerical simulations for $n_s \leq 1$. However, when we increase n_s we see that our model starts to drastically over-predict positively frozen agents as well as under-predict the negatively frozen agents. The major reason for this, from what we think, is the term δ^μ which does not average to 0 in this regime since the assumption of independence of distributions probably is not reasonable. Here there is an obvious possibility to continue the work and look into this closer.

Another obvious thing that we have overseen in our model is the fact that the number of active agents fluctuates in simulations. We disregard this as the first thing we do is taking the time average and apply the time averaged ϕ_i uniformly to each history. This means that we cannot simulate the time evolution of active agents from our model, but we get a good estimate of the average amount of activity. This means that we will probably underestimate the actual fluctuation of the attendance. To account for these fluctuations in active agents is evidently something that could be studied in further works to possibly make the model more accurate. Finally we can mention that there might be a point in developing a more rigorous framework for how to distinguish between fickle and frozen agents.

Appendix A A few analytical calculations

Here we will present some of the calculations of previous chapters that were left out. First we look into how we find the conditional mean of $\vec{a}\cdot\vec{\Omega}$ over the positively frozen agents

$$\overline{(\vec{a}\cdot\vec{\Omega})} = P\overline{(a^\mu\Omega^\mu)} = P\int a^\mu\Omega^\mu \mathrm{P}(a|\Omega)\mathrm{P}(\Omega)\,\mathrm{d}a\,\mathrm{d}\Omega\,.$$

Here we make the variable change $x = a^\mu + g_+\Omega^\mu$ to get zero mean random variables. By doing this we get two terms inside the integral. One is just simply the expectation of x and Ω^μ which will be zero and the other one will be the variance of Ω^μ multiplied with $-g_+$ which is what the integral evaluates to.

The calculation for finding the variance of the mean step length is as follows, again with the assumptions of independence of distributions for large P

$$\sigma_\delta^2 = \frac{1}{P}\int (a^\mu(\Omega^\mu + X^\mu + Y^\mu))^2 \mathrm{P}(\Omega)\mathrm{P}(a)\mathrm{P}(X|\Omega)\mathrm{P}(Y|\Omega)\,\mathrm{d}X\,\mathrm{d}Y\,\mathrm{d}\Omega\,\mathrm{d}a\,.$$

We get that the a^μ term integrates to 1, which is to say it is just the variance of a^μ. We then make variable changes on X^μ and Y^μ in order to get distributions centered around 0. If we do this we get

$$\sigma_\delta^2 = \frac{1}{P}\int (a^\mu(\Omega^\mu(1-N_s(f^+g_+ + kf^0g_0)) + X'^\mu + Y'^\mu))^2 \mathrm{P}(\Omega)\mathrm{P}(a)\mathrm{P}(X'|\Omega)\mathrm{P}(Y'|\Omega)\,\mathrm{d}X\,\mathrm{d}Y\,\mathrm{d}\Omega\,\mathrm{d}a\,,$$

where $X' \sim \mathcal{N}(0,f^+N_s)$ and $Y' \sim \mathcal{N}(0,kf^0N_s)$. Because of independence of variables, and the fact they all are zero mean variables, we only get contributions from the squared terms. This also means we simply get their variances and we end up with

$$\sigma_\delta^2 = n_p(1 - N_s(f^+g_+ + kf^0g_0)^2 + f^+n_s + kf^0n_s.$$

To find the score distributions we need to solve the master equation. For fickle agents we assume the score distribution is time independent. The equation in the continuum limit then becomes

$$\mathrm{P}(x) = \int \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-x'-\Delta_\pm)^2}{2\sigma^2}} \mathrm{P}(x')\,\mathrm{d}x'\,.$$

If we now make the change of variables $x - x' = -s$ and the ansatz that $P(x+s) = P(x)P(s)$ we get

$$\mathrm{P}(x) = \mathrm{P}(x)\int \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(s+\Delta_\pm)^2}{2\sigma^2}} \mathrm{P}(s)\,\mathrm{d}s\,,$$

which means that $\int \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(t+\Delta_\pm)^2}{2\sigma^2}} \mathrm{P}(s) ds = 1$. Multiplying and dividing the integrand with $e^{-\frac{2\Delta_\pm s}{\sigma^2}}$ we get

$$\int \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(s-\Delta_\pm)^2}{2\sigma^2}} e^{-\frac{2\Delta_\pm s}{\sigma^2}} \mathrm{P}(s)\, \mathrm{d}s = 1,$$

and if $\mathrm{P}(s) = e^{\frac{2\Delta_\pm s}{\sigma^2}}$ we are left with an integral over a normal distribution over the entire real line which obviously integrates to 1. We have thus shown the expression for the fickle agents score distribution.

References

[1] W. B. Arthur, Inductive reasoning and bounded rationality, The American economic review 84 (2) (1994) 406–411.

[2] D. Challet, Y.-C. Zhang, Emergence of cooperation and organization in an evolutionary game, arXiv preprint adap-org/9708006.

[3] J. L. Kocevar, Mentor: prof. dr. rudolf podgornik, fmf mentor: doc. dr. sašo polanec, ef ljubljana, 23.2. 2010.

[4] R. Savit, R. Manuca, R. Riolo, Adaptive competition, market efficiency, and phase transitions, Physical Review Letters 82 (10) (1999) 2203.

[5] D. Challet, M. Marsili, Phase transition and symmetry breaking in the minority game, Physical Review E 60 (6) (1999) R6271.

[6] Y. Li, R. Riolo, R. Savit, Evolution in minority games.(i). games with a fixed strategy space, Physica A: Statistical Mechanics and its Applications 276 (1) (2000) 234–264.

[7] W. Zhang, Y. Sun, X. Feng, X. Xiong, Evolutionary minority game with searching behavior, Physica A: Statistical Mechanics and its Applications 436 (2015) 694–706.

[8] H. Kimura, E. Akiyama, Grand canonical minority games with variable strategy spaces, in: Annual Conference of the Japanese Society for Artificial Intelligence, Springer, 2005, pp. 291–301.

[9] D. Challet, M. Marsili, Y.-C. Zhang, Stylized facts of financial markets and market crashes in minority games, Physica A: Statistical Mechanics and its Applications 294 (3) (2001) 514–524.

[10] A. Cavagna, Irrelevance of memory in the minority game, Physical Review E 59 (4) (1999) R3783.

[11] D. Challet, M. Marsili, Relevance of memory in minority games, Physical Review E 62 (2) (2000) 1862.

[12] L. Andersson, S. Bertolino, J. Lindqvist, P. Nilsson Lundberg, S. Rubenson, Finansiella marknader som komplexa system – uppkomsten av empiriska fenomen i en agentbaserad modell av aktiehandeln.

[13] D. Challet, M. Marsili, Criticality and market efficiency in a simple realistic model of the stock market, Physical Review E 68 (3) (2003) 036132.

[14] A. Cavagna, J. P. Garrahan, I. Giardina, D. Sherrington, Thermal model for adaptive competition in a market, Physical Review Letters 83 (21) (1999) 4429.

[15] M. Granath, A. Perez-Diaz, Diffusion and localization of relative strategy scores in the minority game, Journal of Statistical Physics, Volume 165, Issue 1, pp.94-114.

[16] M. Marsili, D. Challet, Continuum time limit and stationary states of the minority game, Physical Review E 64 (5) (2001) 056138.