Learning optimal risk strategies in multi-agent economic systems

Julian Neñer[1], Javier Kreiner[2], and María Fabiana Laguna[3]

[1] Instituto Balseiro, Universidad Nacional de Cuyo, R8402AGP Bariloche, Argentina
julian.nener@ib.edu.ar
[2] javkrei@gmail.com
[3] Centro Atómico Bariloche and CONICET, R8402AGP Bariloche, Argentina
lagunaf@cab.cnea.gov.ar

Abstract. Optimal individual strategies for agents undergoing transactions in the Yard-Sale model were studied. The inclusion of rationality in their behavior by endowing each agent with a neural network and training them with a genetic algorithm showed promising results.

Keywords: Econophysics · Wealth distributions · Machine learning.

1 Introduction

It is a well known fact that many countries around the world display highly uneven wealth distributions[1]. This phenomenon can be observed in all societies to a certain degree, and has been present throughout human history.

Recurring statistical patterns were observed for the first time by the economist Vilfredo Pareto a century ago, when he discovered the existence of power laws in the wealth distributions of european countries[2]. Modern studies have painted a more nuanced picture, in which these power laws are correct only for the highest social classes, while the group of poorer individuals exhibits a log-normal or Gibbs distribution[3]. These observations lead to the following question: is this kind of behaviour inherent to societies themselves?

With the goal of attaining a deeper understanding of the social interactions that lead to this type of wealth distributions, simplified models based on ensembles of economic agents have been proposed. Consider a system of N agents. A pair i and j, with wealths w_i and w_j respectively, is chosen at random, and they take part in a transaction. After their mutual interaction their wealths will evolve according to the following dynamics:

$$\begin{aligned} w_i(t+1) &= w_i(t) + (2\eta_{i,j} - 1)\Delta w_{i,j} \\ w_j(t+1) &= w_j(t) - (2\eta_{i,j} - 1)\Delta w_{i,j}, \end{aligned} \tag{1}$$

where $\eta_{i,j} \in \{0, 1\}$ is a dichotomous random variable, and $\Delta w_{i,j}$ is a quantity determined by the specifics of the model. Note that under such a transformation the total amount of wealth in the system is conserved. One of the most well known models within this family is the so called Yard-Sale model (YSM)[4], where agents interact risking a random fraction of their wealth. The main (and only) interaction parameter is then defined as the risk factor r, which can be different for each agent. The amount of wealth transferred in each interaction will be $\Delta w_{i,j} = \min(r_i w_i, r_j w_j)$. It has been proven that the only possible macroscopic equilibrium for this system is the condensation of all wealth in a single agent[5]. To alleviate this effect, an asymmetry is added to the distribution of $\eta_{i,j}$, that favors the poorest of the agents i and j in a single transaction:

$$p_{i,j} = \frac{1}{2} + f\frac{|w_i - w_j|}{w_i + w_j}, \tag{2}$$

where $f \in [0, \frac{1}{2}]$ is usually called the social protection factor. The macroscopic aspects for this type of models has been thoroughly studied for the past decades[6][7], but studies regarding their microscopic behaviour are lacking.

In this work, we propose a new approach for the study of the YSM in which agents are endowed with rational behaviour in their decisions. Keeping in mind that the only parameter of interaction that agents have is their risk factor r_i, each agent will be allowed to change it at each time step. We are interested in finding the optimal individual strategy. The solution will then be a function $r_i(t+1) = R(\vec{v}_i(t))$ that maps an input vector $\vec{v}_i(t)$ containing all the information available to agent i at time t to the level of risk the agent will use in the following time step.

2 Methods

We use an evolutionary algorithm to evolve a system of N agents, initialized with wealths and risks uniformly distributed in the interval $[0, 1]$. The fitness of each agent is then calculated as the average wealth obtained after T_{gen} Monte Carlo steps (MCS), where a MCS is defined as $N/2$ YSM transactions in the system. A new generation of agents is then created, where each new agent is an imperfect copy of an agent of the previous generation, selected with a probability proportional to its calculated fitness. The system is then reset to a new random initial condition. The process iterates until convergence to a solution is reached. We train for a range of $f > 0$.

We conduct three experiments. First we want to ascertain the existence of an optimal strategy. To that purpose, 1000 systems of 10000 agents each, with fixed risks uniformly distributed in the interval [0,1], are evolved to macroscopic equilibrium, defined as the time at which the Gini index[8] becomes constant. In the next couple experiments, each agent is given an equally structured multilayer perceptron network (MLP) with random weights. The information vector $\vec{v}_i$ is taken as the input at each time step, and the risk of each agent will change according to the corresponding output. First we let $\vec{v}_i(t) = r(t)$, so that the input is the previous step's risk. In the third experiment we let $\vec{v}_i(t) = w(t)$, the wealth in the previous time step.

3 Results

The histograms of the average wealth per agent $\langle w_i \rangle$ as a function of risk were plotted, as shown in the top panel of Figure 1. The presence of local maxima in these curves shows the existence of an optimal risk r_{opt} that increases with f.

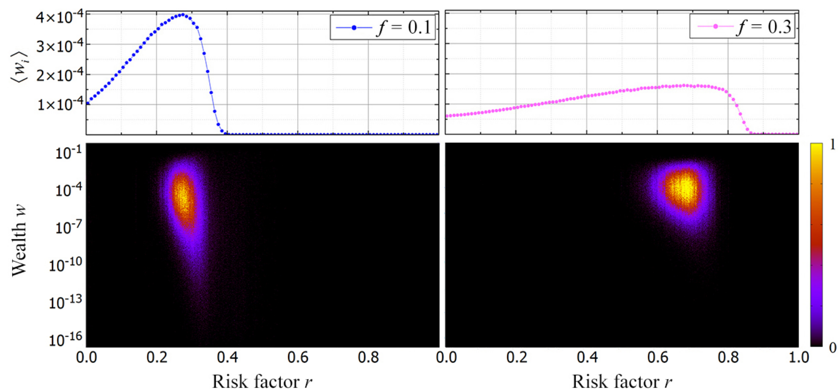

Fig. 1. Risk histograms of the average wealth per agent $\langle w_i \rangle$ for different values of f, belonging to 1000 untrained systems (top), and their comparison with density plots of agents in the $r - w$ plane belonging to 1000 trained systems (bottom).

The simplest form of rational behaviour is introduced in the second experiment, as described in the previous section, by setting $\vec{v}_i = r_i(t)$. That is, every agent can only see their risk at every time step, and change it according to its MLP. It can be shown that this type of function always leads to a constant, optimal r, independent of w. To find the optimal solution, 1000 systems of 1000 agents each were trained in parallel and then brought to equilibrium using CUDA. As the agents are described by two parameters, r and w, a density plot was made in the $r - w$ plane, as shown at the bottom panels in Figure 1. The highest density of agents is centered at the optimal risk r_{opt}, previously found in the histograms, where each agent maximizes their average wealth. This result indicates that the application of rational behaviour through self learning algorithms in this type of systems is possible. It is noteworthy that the Gini index for trained systems always converged to higher values than those of untrained systems, hinting at the existence of higher wealth inequality.

The next step in complexity was giving the agents their wealth w_i as input at each time step. To verify the existence of convergence, 1000 systems were independently trained once again, and then the functions defined by each individual MLP were plotted in a density map. The result obtained for $f = 0.3$ after a reasonably high number of generations ($\sim 10^5$), such that the curve defined by the region of maximum density stops changing, is shown in the left panel of Figure 2. It can be seen that the best solution for high wealth agents is to increase their risk. This behavior can be attributed to the definition of the $\Delta w_{i,j}$ for the YSM (the money exchanged in each transaction): if the wealth of an agent is high, a high risk ensures that the wealth exchanged will always be whatever fraction the other agent is risking, and gives the maximum possible profit for the high-w agent in the next immediate step.

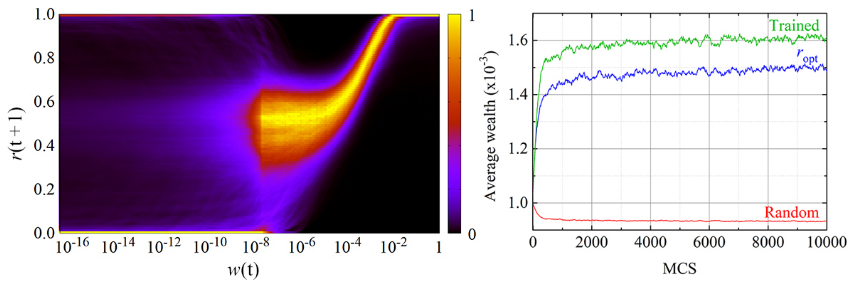

Fig. 2. Left panel: solutions found by the trained agents for $f = 0.3$ after 10^5 generations. Right panel: comparison of the average wealth obtained by 100 trained agents with baselines given by 900 agents with fixed r.

We observe that there is an increasing dispersion as the w values get lower, until no solution is found for $w < 10^{-8}$. This occurs because the agents are not able to reach the region of lower w when they become sufficiently trained, and thus those states are never experienced. To test whether the solution found is successful or not, the average wealth obtained at every time step was compared between the trained agents and two different groups that act as baselines: the group of agents with the previously found optimal risk r_{opt} (fixed), and all the non-trained agents with fixed risks (chosen at random). The result is shown at the right panel of Figure 2, where it can be seen that the trained agents obtain the highest average wealth at all times.

4 Conclusion and Future Work

Microscopic aspects of the wealth distribution model known as the Yard-Sale model were studied, which led to the observation of an optimal strategy for this type of systems. The simplest rational behaviour was then incorporated in the agents by using neural networks trained with a genetic algorithm, and it was then confirmed that every agent could maximize their average wealth.

For future work, the incorporation of additional inputs to each agent's MLP, such as the wealth risked by the opponent, is currently being studied.

References

1. Souma, W.: Universal Structure of the Personal Income Distribution. Fractals: Complex Geometry, Patterns, and Scaling in Nature and Society 9, 463–470 (2001)

2. Pareto, V.: Cours d'Economie Politique. F. Rouge, Lausanne (1897)
3. Pianegonda, S., Iglesias, J. R.: Inequalities of wealth distribution in a conservative economy. Physica A: Statistical Mechanics and its Applications 342, 193–199 (2004)
4. Hayes, B.: Follow the Money. American Scientist 90, 400–405 (2002)
5. Chorro, C.: A simple probabilistic approach of the yard-sale model. Statistics and Probability Letters 112, 35–40 (2016)
6. Chakraborti, A. et al: Econophysics: Empirical facts and agent-based models. Quantitative Finance 11, 991–1012 (2011)
7. Boghosian, B. M.: Kinetics of wealth and the Pareto law. Phys. Rev.E 90, 4 (2014)
8. http://en.wikipedia.org/wiki/Gini_coefficient