Finding optimal strategies in the Yard-Sale model using neuroevolution techniques

Julian Neñer
Balseiro Institute, Argentina

María Fabiana Laguna
Statistical and Interdisciplinary Physics Group,
Centro Atómico Bariloche, Argentina

A new type of in-depth microscopic analysis is presented for the Yard-Sale model, one of the most well known multi-agent market exchange models. This approach led to the classification and study of the individual strategies carried out by the agents undergoing transactions. The findings allowed to determine a region of parameters for which the strategies are successful, and in particular, the existence of an optimal strategy. Strategies that maximize the individual wealth of each agent were then found by incorporating machine learning techniques and performing their training through a genetic algorithm. It was found that the addition of trained agents in these systems leads to an increase in wealth inequality at the collective level.

1. Wealth inequality

It is a well known fact that different countries around the world exhibit highly unequal wealth distributions. This phenomenon has not only been observed in many societies at different scales, but has been present repeatedly throughout history.

These observations lead to the following question: is this kind of behavior inherent to societies themselves?

2. The Yard-Sale model

Simplified models based on ensembles of economic agents have been proposed in the past decades to try to explain this kind of behavior. In particular, the so called Yard-Sale Model (YSM) is one of the most studied in the field.

The YSM is defined in a system of N agents, each of them characterized by their wealth w and a risk propensity r that determines the fraction of wealth that will take part in the interaction.

$$w_i(t+1) = w_i(t) + (2\eta_{i,j} - 1)\Delta w_{i,j} \qquad \eta_{i,j} \text{ r.v.} \in \{0,1\}$$

$$w_j(t+1) = w_j(t) - (2\eta_{i,j} - 1)\Delta w_{i,j} \qquad \Delta w_{i,j} = \min(r_i w_i, r_j w_j)$$

To avoid wealth condensation in a single agent, 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:

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

Simulations made with the YSM are able to replicate certain characteristics found in empirical data, such as wealth distributions that exhibit power-law tails, and exponential behavior at the region of lower income, among others. This model has proven to show interesting results at the macroscopic level, but studies regarding its microscopic aspects have hardly been explored.

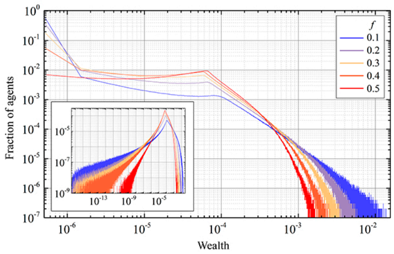

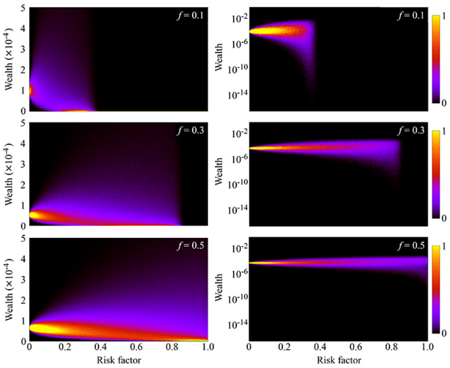

With the goal of studying the individual behavior of the agents and how it leads to the well known macroscopic results, we recorded the final state of a large number of agents (obtained from 10^3 simulations of 10^4 agents each one).

For a better visualization of the data, we define a density function $\rho(r, w)$ which takes increasing values with the amount of agents in a certain region of $r - w$ space. It is of special interest to note the existence of the critical risk factor r_{crit}, such that agents with r above that value will always end up losing all their wealth.

This critical risk allowed for a classification of winner and loser strategies in this model, which depend on the values of the parameters f and r, and also turned out to be scale invariant. Another interpretation of such classification is that it is possible to estimate how much wealth should be risked in different types of societies.

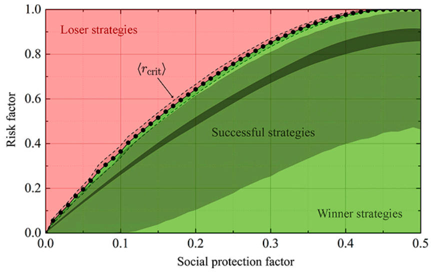

3. Incorporating rationality

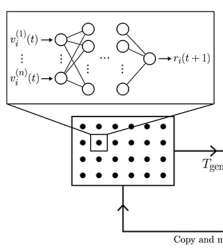

We use an evolutionary algorithm to evolve a system of N agents, initialized with wealths and risks uniformly distributed in the interval (0, 1). The fitness of each agent is then calculated as the average wealth obtained after T_{gen} Monte Carlo steps (MCS), where a MCS is defined as $N/2$ YSM transactions 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.

The histograms of the average wealth per agent $\langle w_i \rangle$ as a function of risk were plotted, as shown in the top panels. The presence of local maxima in these curves shows the existence of an optimal risk r_{opt} that increases with f. Then, systems of $N = 10^5$ rational agents endowed with neural networks with their risks as their only input were trained until reaching convergence.

It is clearly observed that the highest density of agents is centered at the optimal risk previously found in the curves, where each agent maximizes their average wealth.

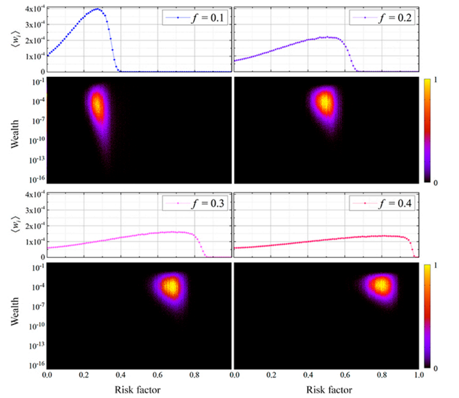

4. The personal wealth as input

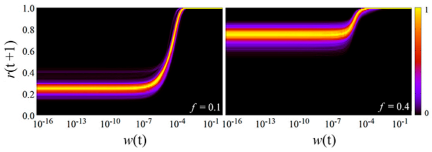

A system with a fraction of 10% rational agents was trained.

When the functions were plotted in a density map after reaching convergence, it was found that the solution consisted in increasing the risk as the wealth increases.

Then, to test whether the solution found is successful or not, the average wealth obtained at every time step was compared between the trained agents, those with r_{opt}, and all of the irrational agents.

It can be seen that the trained agents obtain the highest average wealth at all times.

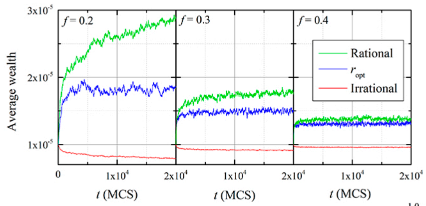

By studying macroscopic aspects of the system such as the Gini index and wealth distributions, a progressive increase in inequality was found when the system had a larger subgroup of rational agents. It was also found that wealth distributions approach power laws in the entire wealth range for certain values of the social protection factor f.

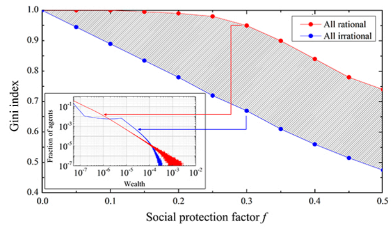

5. Incorporating the opponent's bet

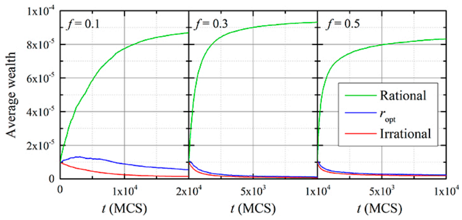

When agents were given both their personal wealth and the wealth fraction risked by the opposing part, it was found that the trained subgroup reaches wealth values an order of magnitude higher than the rest, verifying that the solutions found provide a great advantage. In this case, the solutions found by the agents showed regions in the $w_i - r_j w_j$ plane for which they chose not to interact.

It was also found that they could estimate when it would be convenient to perform a transaction. The trained agents could know when an interaction would be favorable by taking into account the effect of the social protection factor.

They were also capable of being selective by choosing not to interact when their w became sufficiently high.

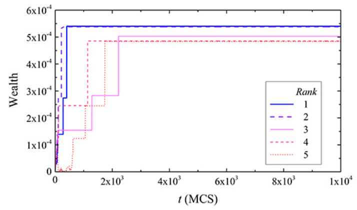

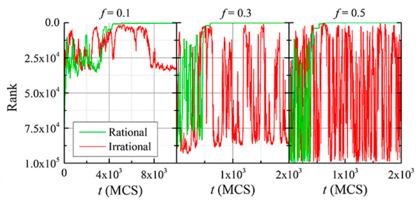

In general, it was found that the fact that every agent is endowed with a certain ambition to maximize their personal wealth, compatible with the individualistic wealth distribution model proposed, always leads to a higher wealth inequality at the collective level.

J. Neñer and M. F. Laguna, Optimal risk in wealth exchange models: Agent dynamics from a microscopic perspective, Physica A **566** (2021).
J. Neñer and M. F. Laguna, Wealth exchange models and machine learning: finding optimal risk strategies in multi-agent economic systems, Physical Review E, in press (2021).