

Games on Cellular Spaces: An Evolutionary Approach

Pedro Ribeiro de Andrade, Antonio Miguel Vieira Monteiro,
and Gilberto Câmara

INPE – National Institute for Space Research,
Av. dos Astronautas 1758, CEP 12227-001, São José dos Campos, Brazil
{pedro,miguel,gilberto}@dpi.inpe.br

Abstract. By using differential equations, evolutionary game theory shows that most of the games of competition for resources have equilibrium strategies named Evolutionary Stable. Although this approach can deduce these points, it is not possible to say how or whether a population will reach such equilibrium. We present an evolutionary agent-based model where individuals compete for space using mixed strategies. Agents belong to spatial locations that settle with whom they can interact, but they can freely move to contiguous partitions according to a definition of satisfiability. The simulation results show that, although the agents do not have any knowledge about equilibrium points, the population's mean strategy always converges to a stable state, close and above to the analytic equilibrium. Moreover, it is reached independently of the initial population.

Keywords: Spatial Games, Agent-Based Modelling, Chicken Game, ESS.

1 Introduction

Since the works of Maynard Smith on evolution and the theory of games [1], it has been shown that Game Theory can be used for studying species competing for resources. By using replicator equations, evolutionary game theory shows that there exist equilibrium points named Evolutionary Stable Strategy (ESS). It means that, if all members of a population adopt this strategy then no mutant strategy could spread under the influence of natural selection.

Although this approach can deduce equilibrium points, it does not say how or whether the population will reach such equilibrium. Computer models are the way to show how this happens by growing the patterns observed in real world, because finding equilibrium points is not enough to show that a population will reach such equilibrium: it is necessary to build these patterns in a generative way [2].

Authors that propose simulation models try to flexibilize some assumptions of the mathematical model, criticizing the fact that some of its assumptions are not applicable to the real world. The critics rely mainly in the use of infinite populations. An example is the work of Orzack and Hines, showing that the probability an ESS will evolve is proportional to the population size [3]. Another example is the work of Fogel and others, which describes models with a population starting at ESS but,

instead of staying stable, the models evolve to other strategies [4]. In both cases, each agent competes with all other agents of the model in order to calculate its fitness. These models have two main disadvantages. First, they rarely converge to a stable state, be it an ESS or not. Second, they are strongly dependant on the composition of their initial population.

By adding space as a new dimension for studying these competitions, we can relax the assumptions of agents having to compete with each other and working with infinite populations. There are some experiments in the literature showing that members of populations compete for space because they need to guarantee food, which can decide their fate. An experimental evidence that owning space may lead to survivor is presented by [5]. The author describes the percentage of diary corporal weight change of spiders. There are two groups of spiders: one that owns some territory and another that does not. Most of spiders that do not own a territory lose weight, while almost all owners increase their weight, which leads to a greater chance to survive.

However, much of the simulation studies that add space as a fundamental component of the model are interested in the evolution of cooperation using only pure strategies or meta-strategies that use the last results to choose the next action [6]. In these works, only the strategy spreads over the space, or, in some other cases, the agents have little mobility, mainly waiting for an empty cell to move [7]. In this work, we study the evolution of a population that compete for space using *mixed* strategies, and whose members can freely move according to a definition of satisfiability.

The main objective of this work is to investigate whether a population evolves to a stable equilibrium within this environment, and how the reached equilibrium is related to the theoretical equilibrium. We present initial results of the proposed model, analysing the effects of mutation and initial population on the evolutive process.

2 Non-cooperative Games, Nash Equilibrium, and ESS

A game with n players is said to be non-cooperative when we have, for each player:

1. a finite set of pure strategies (actions);
2. a payoff function, mapping all n -tuples with pure strategies to real numbers.

One mixed strategy is a collection of non-negative numbers adding up to 1, corresponding to probabilities of using each of the pure strategies. The mixed strategy defines the tendencies of a player. Each time it plays, it will choose randomly one of its pure strategies, based on the probabilities defined by the mixed strategy.

For example, let us take the chicken game. Two players have the choice to escalate (E) or not to escalate (\sim E) a brawl. If none of them escalates, nothing happens. If only one escalates, the other player runs away, and the winner receives 1 from the coward player. But, if both decide to escalate, each player pays 10 due to medical care. This game is said to be symmetric, because both players employ the same pure strategies and payoffs, as shown in Table 1. Given that this game has only two pure strategies, we say that s_x , $0 \leq x \leq 1$, is the mixed strategy of escalating with probability x .

Table 1. Chicken game payoffs, in pairs (A, B)

	B escalates	B does not escalate
A escalates	(-10, -10)	(+1, -1)
A does not escalate	(-1, +1)	(0, 0)

Nash proved that, given any non-cooperative game of n players, there is always an equilibrium point, a set of mixed strategies for each player that, if a player changes its mixed strategy, the best result it may get will be the same as in the equilibrium [8]. No player has incentive to deviate one-sidedly from its strategy as long as the other players remain in the equilibrium. This is known as the Nash equilibrium.

But this concept of equilibrium may cause controversy. Let the chicken game and two players, A and B, following strategies s_a and s_b , respectively. The expected payoff of A is $-10s_a s_b + s_a - s_b$. If A would know exactly the value of s_b , it would be possible to calculate the best action for A. If s_b is greater than 10%, the best choice for A is never to shoot ($s_a = 0$), implying in a payoff of $-s_b$. If s_b is less than 10%, A should always shoot ($s_a = 1$), because its payoff would be $1 - 11s_b$. But, if s_b is exactly 10%, all strategies for A lead to the same payoff (-0.1). Thus, if s_a is also fixed at 10%, no other strategy could increase its payoff against A by changing its own mixed strategy. Applying the same reasoning for B, we arrive to the conclusion that when both players follow $s_{0.1}$ the game is in a Nash equilibrium. Most game theorists agree on $s_{0.1}$ as the rational solution for this game, but the argument is weak [9]. Although deviating from the equilibrium does not increase the utility of a player, it does not decrease as well, as long as the opponent follows equilibrium. Thus, this equilibrium it is not *strict*.

A clearer explanation can be found when it is played not by only two players, but within a population. Maynard Smith viewed this game in a population-dynamical setting. In his model, an infinite number of players meet randomly in contests where they have to decide whether to escalate or not. If the estimated overall probability is greater than 0.1, it is better not to escalate. If it is less than 0.1, it is better to escalate. But if it is exactly 0.1, then there is no better strategy than $s_{0.1}$. In this sense, self-regulation leads to $s_{0.1}$ – self-regulation, not between two players, but within a population. Nash has also proposed a similar interpretation for the equilibrium points, the *mass-action* [8], forgotten for decades in his unpublished thesis.

Maynard Smith pointed out two possible interpretations for the ESS. The first interpretation refers mixed strategies, meaning that each member of the population follows the same equilibrium mixed strategy. In the case of the chicken game, the whole population would follow $s_{0.1}$. The second solution works with pure strategies. We have players following different pure strategies in such a way that the mean strategy of the overall population is at equilibrium. Again with the chicken, the stable state has 90% of the population following $s_{0.0}$ and 10% following $s_{1.0}$.

In this work, we will study how this non-strict equilibrium behaves in a spatial context where the agents' mobility is based on the results of the games. We relax the assumptions of finite population and matches among all members of the population, and add new parameters such as fitness, satisfiability, and mutation.

3 The Basic Model

The approach of this work is based on a non-evolutive model proposed Andrade and others [10]. The model takes place in a cellular space. A cellular space is a network of cells connected by neighbourhood relations. The simplest example of a cellular space is a grid, with square cells having four touching neighbours. The cellular space is populated with agents. Each agent belongs to a cell, which has enough space for it to live. Initially, a cell contains a set of agents, which have to compete for it through a non-cooperative game. Whenever an agent is playing a non-cooperative game, we call it *player*.

The basic assumption of this model is that whenever an agent arrives at a cell it is satisfied with it, and it will not move until it becomes dissatisfied. Two agents within the same cell may play a game competing for it, and the result of the game affects the satisfaction of both agents. This is the only memory an agent has, and it is called the local satisfaction. It starts with a positive value when an agent arrives at a cell, and when this value reaches zero or less, the agent randomly picks a neighbour cell and moves to it, looking for a better cell to compete for. Therefore, this movement is a random walk.

Each agent also has a global satisfaction, starting with a positive value significantly greater than the local satisfaction, but also affected by the payoffs of the games. All agents have the same global satisfaction at the beginning of the model. An agent that got dissatisfied many times and its global satisfaction reaches zero or less leaves the model. As we need local and global satisfaction decreasing along the simulation, the expected payoff of the game used in the model has to be almost always negative.

To create a metric for satisfaction, we say the satisfaction of an agent is measured by its fitness. Local satisfaction represents the maximum effort one can dispend when competing for a cell and global satisfaction is the initial fitness. Agents are identical if we consider satisfaction, but they differ in their mixed strategies, which cannot be identified by any other agent.

The basic model has a finite number of turns, each one with two steps. The first step establishes the games, randomly choosing pairs of agents in each cell, and then carries out the games with each pair. Cells with an odd number of agents have one random idle agent. No agent will play more than once in each turn.

The second step defines the dynamical part of the model. Once each agent already knows its own payoff, it updates its local and global satisfactions with the earned payoff. Then, it checks if any satisfaction has reached zero or less to perform a movement or to leave the model. The model executes until it reaches a stable state, which can be when there is at most one agent in each cell, or when the overall satisfaction stops to decrease.

The model of games on cellular space can be formalized as a 9-tuple:

$$M = (C, n, S, p, A, s, k, g, l), \text{ where}$$

- C is the cellular space in which the games take place,
- n is the number of players involved in the non-cooperative game,
- S is the set of actions (pure strategies) each player can take,
- $p: X \rightarrow \mathfrak{R}, X=$, is the payoff function,
- A is the set of sorts of agents,

- $s: A \times S \rightarrow [0,1], \forall a \in A, \sum_{b \in S} s(a, b) = 1$, represents the mixed strategies,
- $k \in \mathbb{N}$, is initial number of agents of each given sort,
- $g \in \mathfrak{R}$, is the global satisfaction threshold,
- $l \in \mathfrak{R}$, is the local satisfaction threshold.

Therefore, given s_x , we have that $s(s_x, E)=x$ and $s(s_x, \sim E)=1-x$. An agent using a mixed strategy commits to a randomization device. Each time the agent plays, it chooses one of its pure strategies, based on the probabilities specified by its mixed strategy. As an example of model, the following describes the traditional chicken game:

$$M_c = (Cc, 2, u, \{s_a, s_b\}, s, 1, \infty, \infty).$$

The cellular space Cc has one cell with two agents, one following s_a and the other s_b . They have the same set of possible actions, “escalate” and “not escalate,” and the same payoff matrix u (shown in Table 1), but they do not have to follow the same mixed strategy. Both agents always stay in the cell, and never leave the model.

In this work, we apply the chicken within this basic framework. The expected payoff of chicken is almost always negative, only in the case where both players never escalate it is zero. Thus, this game fits in with the need of reducing satisfaction to make the agents move. C is a squared grid with 20×20 cells, such that the possible movements are at most to four neighbours (up, down, left, and right). Cells on the edges have three choices, and cells on the corners have only two.

4 The Evolutionary Model

The evolutive model is based on the idea that agents that own cells at the end of the basic model deserve to pass their strategies to the next generation. Therefore, the model considers a competition of the basic model as one evolutive step, a survival of the fittest. Once a competition finishes, each agent that has conquered a partition of space produces descendants within the cell it has conquered. The agent tries to transfer its own strategy to its offsprings, but their learning ability is limited, and then they may have a strategy slightly different from the father, according to a predefined mutation probability. Whenever this mutation is activated, the offspring has the strategy of its father with some random change; otherwise it takes exactly the same strategy. Therefore, there are three additional parameters if compared with the basic model: number of descendants, mutation probability, and the mutation itself.

The father leaves the model right after it generates its descendants. This leads to non-overlapping generations, which is a common approach in the literature. This evolutive process can be repeated indefinitely, and it depends only on the initial population of the basic model.

The model has five stochastic ingredients, three of them coming from the original model, and the other two added by the evolutive part. They are:

1. selecting pairs of agents within the same cell to play,
2. choosing the agent’s pure strategy based on its mixed strategy,
3. relocating, once one agent is unsatisfied,
4. mutating or not a given strategy to a descendant, and
5. choosing the mutation change, once it is activated.

5 Experiments

In this work, we present early results that examine the convergence of a population and the effect of mutation on evolution. The parameters chosen were: 3 as the number of descendants, 0.1% as chance of mutation, and ± 0.1 as the change in the inherited strategy, with 50% of probability of each, once the mutation is activated. Agents following $s_{0,0}$ and $s_{1,0}$ can only produce mutants of $s_{0,1}$ and $s_{0,9}$, respectively.

The model was implemented using TerraME framework [11]. Although the model has five random ingredients, the simulations had similar development. Thus, the results presented in this section are single runs of the model.

5.1 Evolutionary Equilibrium

As initial test for studying the evolutionary capabilities of the model, we examine a model whose population is initially composed by $s_{1,0}$ agents. The model has the following arrangement:

$$M_0 = (Cc, 2, \{E, \sim E\}, u, \{s_{1,0}\}, s, 1200, 200, 20).$$

Figure 1 shows the results, with the left side showing a zoom in the first generations of the simulation. We can see that $s_{1,0}$ starts filling the whole cellular space, but it cannot maintain this situation, stabilizing with agents in 75% of the cellular space. This happens because two agents within a cell may shoot, reach the threshold, and leave the cell.

When the first mutation to $s_{0,9}$ arises in the population, this new strategy rapidly spreads, crushing $s_{1,0}$ and dominating the whole cellular space. This strategy can then produce agents following $s_{0,8}$, then $s_{0,7}$, and so on, each one surpassing the previous strategy, until $s_{0,1}$ appears, as shown in the left part of Figure 1. But $s_{0,1}$ cannot get rid of $s_{0,2}$ and both compete indefinitely. The results agree with [10], that state that as more frequently an agent escalates, greater is its ability to realize cells with a higher




Fig. 1. Development of strategies in a simulation with population of $s_{1,0}$

number of agents that also escalate frequently, as it will reach its threshold for moving faster than the ones that escalate less. Moreover, it gets even more important because the agents have limited fitness and need to stay alive. This leads to an increase in the surviving chance of sub-optimal strategies such as $s_{0.2}$ and $s_{0.3}$. Therefore, although there is a single theoretical equilibrium, there is no single best strategy in this competition for space.

Although neither strategy can surpass all other strategies and dominate the overall population, an evolutive process indeed occurs. If we calculate the average strategy of the population by summing the mixed strategies of every owner at the end of a generation and divide by the total number of agents, we can verify that the population indeed converges to a stable state. In Figure 2, we have the average strategy of the population, with a dashed line showing the theoretical equilibrium point. The average strategy converges to a stable state before the 200th generation and stays close and above the equilibrium point until the end of the simulation.

The generative result of this model is a mixing of the two definitions of equilibrium pointed by Maynard Smith. In the results, we do not have the overall population following the same equilibrium strategy, nor agents following only pure strategies, but what really happens is a combination of mixed strategies whose average is a stable equilibrium point.




Fig. 2. Convergence of the mean strategy using different initial populations

5.2 Mutation Effects

A question that arises from the result of the initial model is how the mutation rate can affect the model development. Eight simulations were carried out, each one with a single mutation rate, varying from 0.1% (the previous experiment) to 50%. It is not feasible to have the mutation rate greater than 50% because, in this case, most of the descendants would follow a strategy different from their predecessors.

It is straightforward that increasing mutation leads to a faster convergence to a stable state, as shown in Figure 3. With mutation of 0.1%, it takes more than 150 generations to converge, while with 50% it converges in less than 25 generations. But note

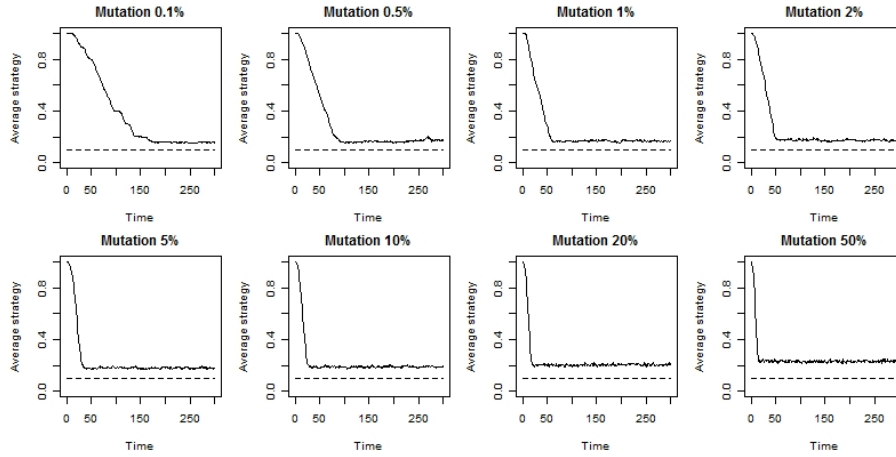


Fig. 3. Effects of changing the mutation rate in the convergence of the model

that the distance between the average strategy and the theoretical equilibrium is proportional to the mutation rate. It is difficult to see comparing two consecutive mutation rates, but we can see clearly comparing the results of 0.1% and 50%.

Mutation induces other impacts in the development of the model. Increasing it leads to a greater diversity of strategies, and then to a higher probability of generating far from equilibrium strategies. However, there is a consequence of increasing the mutation rate that is not straightforward. Figure 4 shows the number of agents following each strategy at the end of each generation, in the same simulations of Figure 3. After the model converges to a stable mean strategy, each individual strategy reduces its oscillation as the mutation increases, stabilizing in a well-defined small interval.

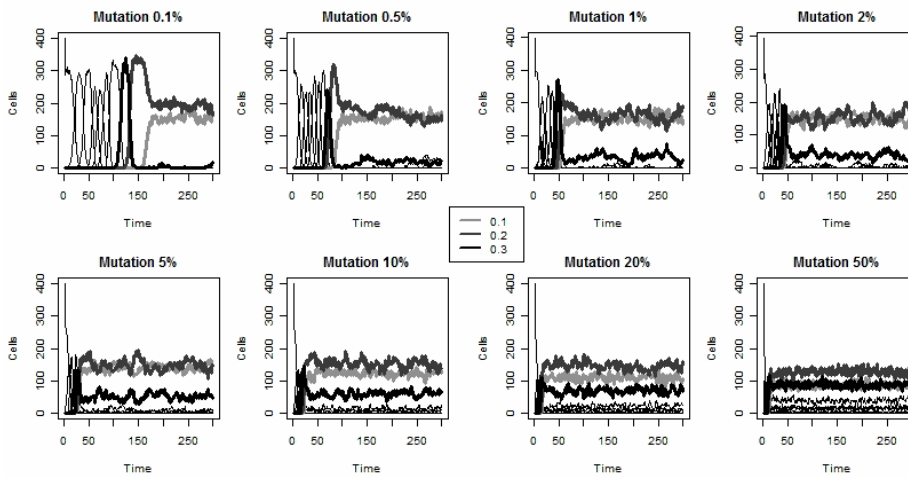


Fig. 4. Effects of mutation on the stability of the final distribution of individual strategies

That happens because, with higher mutation, the model has a more homogeneous distribution of strategies in space, being difficult for a strategy to conquer many cells, and easy for it to maintain a minimum number of cells.

The increase in the mutation rate makes possible the appearance of agents following almost each possible strategy. Strategies with a higher chance of escalating have a major disadvantage when they interact with other agents following a similar strategy. But, as they can avoid themselves moving to other cells, it is possible to keep a few of them in the model. Consequently, these few agents away from equilibrium increase the average strategy significantly, as shown in Figure 3.

The increase in the mutation rate also makes the distribution of individual strategies converge towards a discretized gamma distribution. This result can be seen in Figure 5, where $s_{0.1}$ decreases and $s_{0.3}$ increases as mutation increases, until they almost draw when we have mutation of 50%. Note that increasing mutation allows the development of strategies away from the theoretical equilibrium, making the overall distribution smoother.

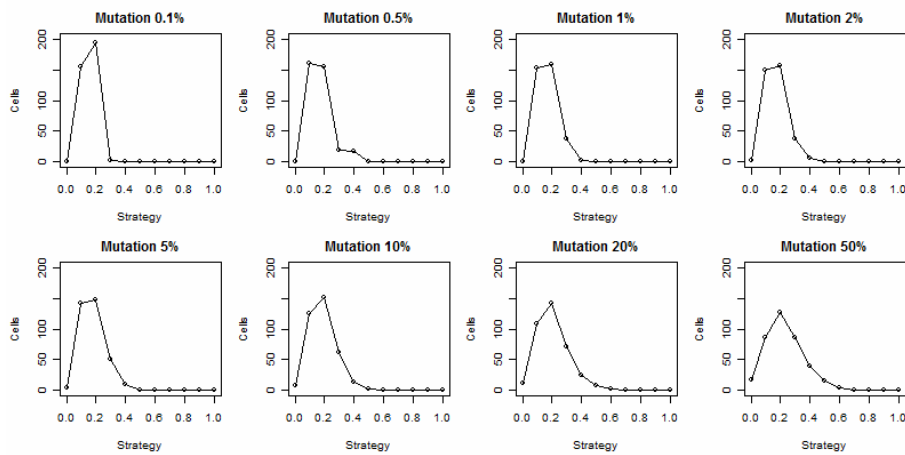


Fig. 5. Mean number of agents of each strategy after convergence, with eight mutation rates

Another result of changing the mutation rate is the number of agents that survive at the end of each simulation. There can exist some empty cells at the end of the simulation because, when there are only two agents within a cell, both might escalate, reach the threshold, and decide to leave the cell. Figure 6 shows the number of agents that give rise to descendants in each generation. The eight simulations have a minimum around 70% of the cellular space in the early generations. After that, the number of agents grows and stabilizes. In each graphic, there is a dashed line pointing out 90% of occupation of the space. In mutation 0.1%, the population stabilizes below this line, while in mutation 10% or more, the mean number of cells surpasses the line. A greater diversity in the number of strategies allows a better occupation of the cellular space at the end of each generation. As there are more agents with different strategies with the increase in the mutation, the fitness of the agents is reduced more heterogeneously, which leads to a lower probability of two agents remove themselves from the model simultaneously.

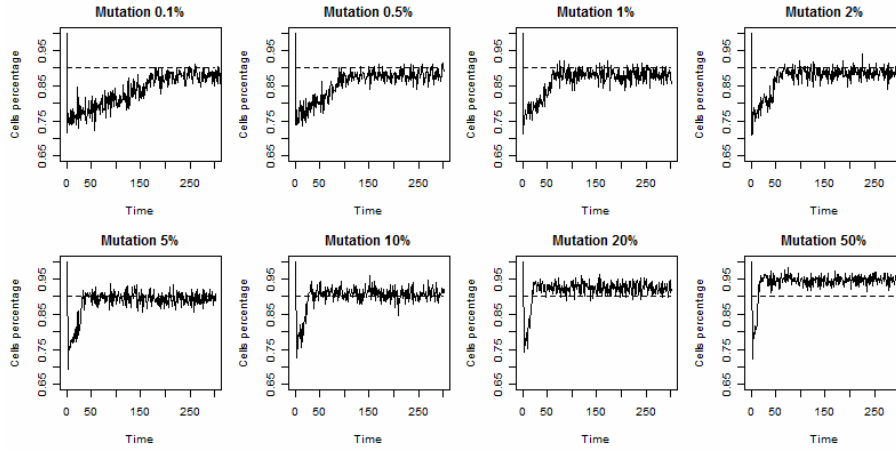


Fig. 6. Effects of mutation on the number of occupied cells at the end of each generation. Note that it also means the number of agents, once we have one agent in each occupied cell.

5.3 Different Initial Populations

The previous results show the impact of mutation over the evolution of the population, but they cannot enlighten whether the evolutive process always happens. To test the evolutionary capability of the model, we need to examine different initial populations. This section presents four models, each one with a single initial population, with the following arrangements:

$$M_k = (Cc, 2, \{E, \sim E\}, u, \{s_k\}, s, 1200, 200, 20),$$

where $k \in \{0.0, 0.1, 0.5, 1.0\}$. That is, the first population of each model has a single strategy. They are: never escalate, equilibrium, random, and always escalate.

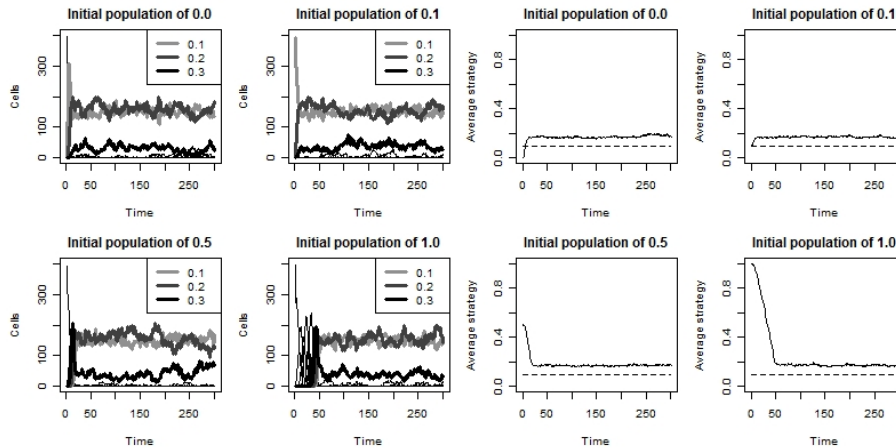


Fig. 7. Individual strategies using different initial populations

Figure 7 shows the results of the four models. In the left, it contains the number of agents following each strategy at the end of each generation, with the three most successful strategies ($s_{0,1}$, $s_{0,2}$, and $s_{0,3}$) as coloured lines. These strategies have similar development in the four cases, with none of them surpassing the other two and dominating the whole population, similarly to the previous results. In the right side, the figure shows the mean strategy of the overall population in each simulation. The four simulations converge to the same stable state, despite the initial arrangement of the models. Both results show that the model is capable to converge to a stable state independently of the initial population. It makes the model stronger, and shows that, although there are five random ingredients in the model, their effects do not make the model chaotic.

6 Concluding Remarks

In this work, the evolution of a population is defined by adding simple rules of reproduction and mutation over a previous work in the literature. The selective step, proposed by Andrade and others, selects the best strategies in a given population, whatever average strategy it follows. The mutation allows some change in the strategy when creating a new generation, which can represent an error when passing the strategy to the descendants. This error leads to the diversity on the population and allows the evolutive process to take place.

The results of our experiments show that, although the agents do not have any knowledge about equilibrium points and cannot even change their individual strategies along their lifetime, the average strategy of the population always converges to a stable state, close and above to the analytic equilibrium. Moreover, this equilibrium is reached independently of the initial population. These two points corroborate the hypothesis that populations evolve to a stable state even if we use a finite population. The results also agree with the statements of Nash, who said that “we can only expect some sort of approximate equilibrium, since [...] the stability of the average frequencies will be imperfect” [8]. As it is almost impossible to validate this kind of model with real world data, showing that the model is independent of the initial population is a reasonable way to make it trustworthy.

A couple of questions are still not solved within the context of the proposed model. As changing the chance of mutation produce effects on the evolutive process, other parameters may also lead to different results. We can cite:

- Reduce the interval of mutation, for example from ± 0.1 to ± 0.01 , or even allow the complete $[0,1]$ range,
- Change parameters such as initial population size and fitness, number of descendants, and mobility threshold, and
- Use games with other equilibrium points.

If we change these parameters, will the model converge to closer and above to the theoretical equilibrium point? Another point that worth some investigation is, can these results be inferred from mathematical equations, or only simulation models can produce that?

References

1. Maynard Smith, J.: *Evolution and the Theory of Games*. Cambridge University Press, Cambridge (1982)
2. Epstein, J.M., Hammond, R.A.: Non-Explanatory Equilibria: An Extremely Simple Game With (Mostly) Unattainable Fixed Points. *Complexity* 7(4), 18–22 (2002)
3. Orzack, S.H., Hines, W.G.S.: The Evolution of Strategy Variation: Will an ESS Evolve? *Evolution* 59(6), 1183–1193 (2005)
4. Fogel, G.B., Andrews, P.C., Fogel, D.B.: On the instability of evolutionary stable strategies in small populations. *Ecological Modelling* 109, 283–294 (1998)
5. Riechert, S.E.: The consequences of being territorial: spiders, a case of study. *American Naturalist* 117, 871–892 (1981)
6. Nowak, M.A., Sigmund, K.: Games on Grids. In: Dieckmann, U., Law, R., Metz, J.A.J. (eds.) *The Geometry of Ecological Interactions: Simplifying Spatial Complexity*, pp. 135–150. Cambridge University Press, Cambridge (2000)
7. Epstein, J.M.: Zones of Cooperation in Demographic Prisoner’s Dilemma. *Complexity* 4(2), 36–48 (1997)
8. Nash, J.: *Non-cooperative games*. In: Mathematics Department, Princeton University Press, Princeton (1950)
9. Sigmund, K.: *Games of Life*. Oxford University Press, Oxford (1993)
10. Andrade, P.R., et al.: Games on Cellular Spaces: How Mobility Affects Equilibrium. *Journal of Artificial Societies and Social Simulation* 12(1), 5 (2009)
11. Carneiro, T.G.S.: Nested-CA: a foundation for multiscale modeling of land use and land change. In: Image Processing Division. INPE, São José dos Campos (2006)