

New insights on the Emergence of Classes Model

David J. Poza García¹, Félix A. Villafañez Cardeñoso², Javier Pajares Gutiérrez³
Adolfo López-Paredes⁴ Cesáreo Hernández Iglesias⁵

Grupo InSiSoc. ETSII Valladolid. Universidad de Valladolid. Pso. del Cauce S/N, 47011
Valladolid.

{¹djpoza@gmail.com, ²villafafelix@yahoo.es, ³pajares@insisoc.org, ⁴adolfo@insisoc.org,
⁵cesareo@insisoc.org}

Abstract. There is an increasing interest in the Social Simulation community in the complete characterization of Agent-Based Models and replication is maybe the central issue. In this paper we propose to extend the replication task and analyze the validity of the results obtained in those models. In many published works authors don't show if obtained results are general or a particular case that depends on particular initial conditions. We present the replication of a well-known published model (Axtell et al. 2000) and we extend the initial work studying how the payoff matrix and the rules initially proposed affect the results and the main conclusions of the original paper.

Keywords: Replication, Agent-based modeling, Validation, Verification.

1 Introduction

If both social sciences and economics are experimental sciences, they need a laboratory (López-Paredes et al., 2002). To this aim, agent based modeling has become an extremely useful methodology, as involving humans in experiments is always difficult, because of restrictions in time and availability (among others).

By means of bottom-up models, social scientists have been able to analyze emergent social phenomena beyond the traditional simulation and experimental techniques. From micro behaviours and interactions among agents, we have been able to build stylized models explaining some of the relevant macro-observed facts.

The pioneering works by Schelling and Axelrod showed us how computational sciences might help social scientists to develop models based on real assumptions about the behaviour of economic agents. However, we had to wait for several decades to put those pioneering works to the test. Thus, the research by Schelling (1971) was intensively extended by Epstein & Axtell (1996), who built up a real Universe (Sugarscape) by means of simple rules. More recently, Galan & Izquierdo (2005) discussed the meta-norm models by Axelrod (1986).

The efforts for replicating previous published models have grown during recent years. However, model replicating is a very tough task, as it was showed by Axelrod et al. (1997), Edmonds & Hales (2003) and Sansores & Pavón (2005). More recently, Wilensy & Rand (2007) proposed some interesting recommendations for improving

diffusion and rigour in multi-agent simulations. Anyway, replication is always the first step for improving and extending previous models, so that new hypothesis and new agent behaviours could be tested.

In this paper, we replicate the model by Axtell et al. (2000) (hereafter AEY), where two agents want a portion of the same pie, and the portion a particular agent gets depends on the portion demanded by the other agent. Our results are in agreement with their conclusions, both with non distinguishable and distinguishable agents (the tag model), as Dessalles et al. (2007) also confirmed in a previous replication of this work. However, in this paper, we analyze the hypothesis that researchers should make to obtain the results shown in AEY'S model. We will pay special attention to the initial conditions of the system (i.e., memory of the players before the game is started) and different ways in which an agent can take a decision. These considerations should be carefully explained in publications to facilitate replication and prevent researchers from making erroneous hypothesis and considering particular cases as general conclusions.

But we try to go one step further. First, we have considered possible artefacts/biases (Galán et al., 2009 & Kubera et al., 2009) and we have tested the results to minor changes in the agents decision rule (as López-Paredes et al., 2004 suggested), so that their decision depends on the most likely option taken by their opponents in previous games; in particular, agents decide based on the opponents decision "statistical mode". It is consistent with experimental research done in neuroscience which demonstrates that humans don't use statistical properties in their internal decision processes.

Secondly, we have tested how dependent the results are on the reward values in the payoff matrix, to see how it affects the aggregated observed behaviour.

The main result of our research is that these simple changes may affect dramatically how and when the equilibrium is reached. Our results confirm the important role of tags in the evolution of the system, what has been empirically demonstrated by Ito et al. (2007) that plays a main role in 'rational' decisions.

2 The Model

We begin by replicating the bargaining model by AEY in which two players demand some portion of a pie. They could demand three possible portions: low, medium and high. As long as the sum of the two demands is not more than 100 percent of the pie, each player gets what he demands; otherwise each one gets nothing.

The authors assume a population of n agents that are randomly paired to play. Each agent has a memory in which he retains the decision taken by his opponents in previous games. The agent uses the information stored in his memory to demand the portion of the pie that maximizes his/her benefit (with probability $1-\epsilon$) and randomly (with probably ϵ).

At first, the authors assume that agents are indistinguishable from one another, but from their memories about previous games. They conclude that, whenever there are not observable differences among the agents (the agents have not a distinguishable tag), there is only one possible state of equilibrium in which all the agents demand

half of the pie. Otherwise, all the agents are either aggressive or passive (some of them demand low and some of them demand high), and no equilibrium is reached.

Secondly, the authors let the agents be distinguishable from one another by introducing a tag: they create two types of agents, each of whom with a different tag. The agents are capable of identifying their opponents' tag and they keep the portion of the pie demanded by their opponents in their memories, both with the same and different tag. In this case, the authors prove that, just by adding different tags to the players, discriminatory states can emerge under certain conditions, in which agents with different tags follow different behaviours.

3 The Model with one agent type

3.1 Replication

First, we have replicated the AEY's model. We used the original payoff matrix (i.e. the combination of values for the different demands): 30 percent for low; 50 percent for medium and 70 percent for high. We also used the original decision rule.

This means that when two players are paired to play, each one gets the portion that they demand as long as the sum of the two demands is less than the 100 percent of the pie. For example:

- if player 1 demands 30 (low), he will receive 30 independently of what player 2 demands.¹
- if player 1 demands 50 (medium), he will get 50 unless player 2 demands 70.²
- if player 1 demands 70 (high), he will get 70 only if player 2 demands 30.³

Problem approach

Using mathematical notation, the payoff matrix can be explained as follows:

n - number of agents

ε - uncertainty parameter

m - memory length of each agent

S_i - space of agent i ($i = 1, \dots, n$) possible strategies

j - possible strategy $\Rightarrow j \in [L, M, H] / M = 50, H = 100 - L, L < H$

(L - select Low, M - select Medium, H - select High)

$[v_1, v_2, \dots, v_m]^i$ - memory array of agent i , which stores the strategies $v_k \in [L, M, H]$ chosen by the opponents in the m previous rounds

¹ When player 1 chooses 30, the sum of 30 (player 1's demand) and all the possible combinations of demands for player 2 are less or equal than 100 percent of the pie.

² If player 2 chooses 70, the sum of the two demands is higher than 100 percent of the pie. In this case, both players get nothing.

³ If player 2 chooses 50 or 70, the sum of the two demands exceeds 100 percent of the pie and each agent gets nothing.

$[A, B]$ - couple of agent randomly paired ($n/2$ randomly pairs by round).

If agent A chooses strategy $i \in S_A$, and agent B chooses strategy $j \in S_B$, they will receive $[i, j]$ if $(i + j) \leq 100$, and $[0, 0]$ if $(i + j) > 100$ (see Table 1, Combination of payoffs)

Decision rule:

What makes an agent choose among low, medium or high? An agent will check his memory to find how often each option has been chosen by his opponents. Then, the agent considers that the probability that his current opponent chooses 30 (low) – for example – is equal to the relative appearance of 30 in his memory. In the same way, he calculates how likely it is for the opponent to choose 50 and 70. Once the agent knows this information, the agent estimates his benefit for the three possible options as follows:

- The average benefit I get if I choose 30 is 30 multiplied by the probability that my opponent chooses 30, 50 or 70.⁴
- The average benefit I get if I choose 50 is 50 multiplied by the probability that my opponent chooses 30 or 50.⁵
- The average benefit I get if I choose 70 is 70 multiplied by the probability that my opponent chooses 30.⁶

Notice that this ‘rational behaviour’ takes place with probability $1-\varepsilon$. However, a random decision is taken with probability ε .

This decision rule is explained with mathematical notation below:

n_j^A - number of positions with value $j \in [L, M, H]$ in the memory array of agent A
 $\Rightarrow [v_1, v_2, \dots, v_m]^A$

$Pr(B_j^A) = n_j^A / m \Rightarrow$ Probability estimated by the agent A for the possibility that the opponent B selects the strategy j (equivalent to the relative frequency of occurrence of value j in the memory array of the agent A)

The utility function for agent A when selects the strategy $i \in S_i = [L, M, H]$ is:

$$U(A_i) = i \cdot \sum_{j \in S_B} [Pr(B_j^A) \cdot V(i, j)] \quad / \quad i \in S_A; \quad V(i, j) = 1 \text{ if } (i + j) \leq 100; \\ V(i, j) = 0 \text{ if } (i + j) > 100$$

Then, each agent A selects with probability $(1-\varepsilon)$ the strategy i that maximizes its utility function:

$$A \text{ select } i \in S_A = [L, M, H] \quad / \quad EU(A_i) = \max U(A_i)$$

And selects a random strategy $i \in S_A$ with probability ε .

⁴ He assumes that the probability that his opponent chooses 30, 50 or 70 is the sum of elements equals to 30, 50 and 70 in his memory divided into the memory size.

⁵ He assumes that the probability that his opponent chooses 30 or 50 is the sum of elements equals to 30 and 50 in his memory divided into the memory size.

⁶ He assumes that the probability that his opponent chooses 30 is the sum of elements equals to 30 in his memory divided into the memory size.

Example

$n = 10; m = 5;$

$L = 30, M = 50, H = 70$

$\Rightarrow S_A = [L, M, H] = [30, 50, 70]$ - space of possible strategies for agent A

if $[v_1, v_2, \dots, v_m]^A = [30, 30, 50, 70, 30]$ - current memory array of agent A

$\Rightarrow n_{30}^A = 3, n_{50}^A = 1, n_{70}^A = 1 \Rightarrow Pr(B_{30}^A) = \frac{3}{5}, Pr(B_{50}^A) = \frac{1}{5}, Pr(B_{70}^A) = \frac{1}{5}$

$$U(A_{30}) = 30 \cdot Pr(B_{30}^A) \cdot V(30, 30) + 30 \cdot Pr(B_{50}^A) \cdot V(30, 50) + 30 \cdot Pr(B_{70}^A) \cdot V(30, 70) = 30 \cdot \frac{3}{5} \cdot 1 + 30 \cdot \frac{1}{5} \cdot 1 + 30 \cdot \frac{1}{5} \cdot 1 = 30$$

$$U(A_{50}) = 50 \cdot Pr(B_{30}^A) \cdot V(50, 30) + 50 \cdot Pr(B_{50}^A) \cdot V(50, 50) + 50 \cdot Pr(B_{70}^A) \cdot V(50, 70) = 50 \cdot \frac{3}{5} \cdot 1 + 50 \cdot \frac{1}{5} \cdot 1 + 50 \cdot \frac{1}{5} \cdot 0 = 40$$

$$U(A_{70}) = 70 \cdot Pr(B_{30}^A) \cdot V(70, 30) + 70 \cdot Pr(B_{50}^A) \cdot V(70, 50) + 70 \cdot Pr(B_{70}^A) \cdot V(70, 70) = 70 \cdot \frac{3}{5} \cdot 1 + 70 \cdot \frac{1}{5} \cdot 0 + 70 \cdot \frac{1}{5} \cdot 0 = 42$$

Agent A selects 70 with probability $(1-\epsilon)$, as it maximizes its utility function.

$EU(A_{70}) = \max U(A_i) = 42$

And selects an random strategy $i \in S_A = [30, 50, 70]$ with probability ϵ .

A simulation of this replication is shown in Figures 1 and 2. Both simulations were run with the same initial parameters (the same number of agents, the same memory size and the same uncertainty parameter $-\epsilon$).

This simplex represents the memory state of the agents. The more demands of L an agent keeps in his memory, the closer to the bottom-right vertex it is plotted. Equivalently, if a player's memory contains a considerable amount of H's, it is placed near the top vertex. Finally, if most of the elements in an agent's memory are M's, it is plotted close to the bottom-left vertex.

The triangle is split into three different regions, separated by three 'decision borders'. The top region is dominated by frequent demands of H in previous matches. This is why agents in this region tend to demand L⁷, as it maximizes their estimated benefit. On the right region, agents are likely to demand H⁷ because L is the dominant element in their memories. Agents on the left region have often found that their opponents demand M; because demanding M maximizes the expected payoff, they are likely to choose M⁷ in the current iteration.

The three 'decision borders' intersect in a point that represents Nash's equilibrium in which agents have the same preference for L, M or H.

AEY states that the system reaches an 'equitable equilibrium' when all the agents have, at least, $(1-\epsilon) \cdot m^8$ elements in their memories equal to M. Figure 1 shows an equitable equilibrium. In this state, all the agents have found frequent demands of M in the past, and they assume that M is the best response. Because all the agents demand M, all the pie is shared out among the players, which means that the system has reached an efficient state. Once the equitable equilibrium is established, it is very

⁷ With probability $1-\epsilon$

⁸ Where ϵ is the uncertainty factor and m is the memory size

difficult for the system to leave this state: the noise parameter ε makes agents choose random probabilities at times. However, this individual behaviour is not able to affect the inertia of the system. Figure 2, by contrast, shows a fractious state, in which all the agents are aggressive or passive (most of them select L or H; M is hardly chosen) and no equilibrium is reached. In this case, the system was started with different random initial conditions. Because agents have not learnt to compromise, some portions of the pie remains undistributed, which shows the high inefficiency of this system.




Fig. 1. Replication of AEY's model with a number of agents $n=100$, uncertainty parameter $\varepsilon=0.2$ and memory size $m=30$. Equitable equilibrium




Fig. 2. Replication of AEY's model with a number of agents $n=100$, uncertainty parameter $\varepsilon=0.2$ and memory size $m=30$. Fractious state.

AEY studied the system behaviour for several values of memory size (m) and uncertainty factor (ε), and so we did in our replication. Figure 3 represents the transition time between regimes (i.e. the number of iterations that it takes for the system to reach an equitable equilibrium starting from a fractious state).

Both models, AEY's and our replication, produce the same result in relation with the transition time: it increases as the memory size grows. Notice that this simulation starts in a fractious state⁹; this is why, at first, all the agents tend to demand L or H with high probability ($1-\varepsilon$) because their memories contains 30 and 70 values. This situation provokes that the agents continue demanding L or H (M never maximizes their expected benefit⁹). Therefore, we depend on the noise parameter ε to escape the fractious state, because this is the only way to make M appear in the agent's memories, and, consequently, make the agents consider M as a good option. When the

⁹ When the system is in fractious state.

system is started (fractional state), the probabilities that an agent chooses M is $\varepsilon/3^{10}$. This is the reason why the higher ε , the higher the probabilities of leaving the fractionous state and thus, the faster the convergence to an equitable equilibrium, as figure 3 shows.




Fig. 3. Replication of AEY's model.. Number of iterations to equitable equilibrium, as a function of the memory size; $n=10$; various ε (uncertainty factor).

3.2 Introduction of a new decision rule

After replicating the original scenario, we changed AEY's decision rule so that the agents demanded the pie portion maximizing their benefits against the most likely option taken by their opponents in previous games (mode of their memory). In this case, an agent assumes that his opponent's option will be the mode of the content of his memory.

An agent will choose H if L is the most frequent decision taken by his opponents in the previous matches; If the most repeated value in his memory is M, the player will choose M. If previous matches show that H is the most frequent decision taken by his opponents, the agent will choose L.

The new decision rule is mathematically explained below:

Each agent A selects, with probability $(1 - \varepsilon)$, its strategy i according to the statistical mode (Mo) of its memory array as follows:

$$Mo[v_1, v_2, \dots, v_m]^A = i / \max n_j^A = n_i^A \text{ for all } j \in S_A = [L, M, H]$$

If $Mo[v_1, v_2, \dots, v_m]^A = L \Rightarrow A \text{ selects strategy } i=H$

¹⁰ The probability of taking a random decision is equal to ε . Supposing that this is the case, the probability that the random decision is equal to M is one out of three (chances that M and H are not randomly chosen). In conclusion, the probability that M is chosen is $\varepsilon/3$.

*If $Mo[v_1, v_2, \dots, v_m]^A = M \Rightarrow A$ selects strategy $i=M$
 If $Mo[v_1, v_2, \dots, v_m]^A = H \Rightarrow A$ selects strategy $i=L$
 and selects a random strategy $i \in A$ with probability ε .*

Example

$n = 10; m = 5;$

$L = 30, M = 50, H = 70$

$\Rightarrow S_A = [L, M, H] = [30, 50, 70]$ - space of possible strategies for agent A

if $[v_1, v_2, \dots, v_m]^A = [30, 30, 50, 70, 30]$ - current memory array of agent A

$\Rightarrow n_{30}^A = 3, n_{50}^A = 1, n_{70}^A = 1 \Rightarrow Mo[30, 30, 50, 70, 30] = 30 \Rightarrow$ Agent A selects 70

with probability $(1 - \varepsilon)$, and selects a random strategy $i \in S_A = [30, 50, 70]$

with probability ε .

When the agents used this new decision rule, the chances of reaching the equity equilibrium were considerably reduced (as López-Paredes et al., 2004 concluded). Figures 4 and 5 show this comparison. To perform this simulation, all the agents were initialized with random memories (as they were in AEY's model), and we measured the percentage of experiments that reached an equitable equilibrium, versus the number of experiments that got stuck in a fractious state.




Fig. 4. Replication of AEY's model. Percentage of experiments reaching an equitable equilibrium. Uncertainty parameter $\varepsilon=0.2$. Original decision rule.




Fig. 5. Replication of AEY's model. Percentage of experiments reaching an equitable equilibrium. Uncertainty parameter $\varepsilon=0.2$. New decision rule.

Figure 4 shows that, provided n (number of agents) and m (memory size) large enough, the probability of reaching an equitable equilibrium is close to 1. However, when the decision rule is changed, the probability of reaching the equity equilibrium and the probability of reaching a fractious state are quite similar, as figure 5 shows.

Furthermore, even when the equity equilibrium was reached, the time to get it was longer in comparison with the same conditions in the experiment with AEY's decision rule.

Figure 6 shows two simulations of our modification of AEY's model, in which the decision rule has been changed as described before. The left triangle shows an equitable equilibrium and the one on the right displays a fractious state. The simulation was run with the same parameters in figures 1 and 2 (100 agents, memory length = 30 and $\varepsilon = 0.2$). Notice how the 'decision borders' change after introducing the new decision rule.




Fig. 6. Modification of AEY's model with a new decision rule. Number of agents $n=100$, uncertainty parameter $\varepsilon=0.2$ and memory size $m=30$. Equitable equilibrium and fractious state.

3.3 Introduction of a variable payoff matrix

In AEY's model, the values of the possible demands are fixed: 30 percent of the pie for low; 50 percent of the pie for medium and 70 percent of the pie for high. We have studied different combinations for low (L) and high (H) rewards to analyse the effects on the behaviour of the system¹¹. The combination of payoffs is shown in Table 1.

P1\P2	H	M	L	P1\P2	H	M	L	P1\P2	H	M	L
H	0 \ 0	0 \ 0	95 \ 5	H	0 \ 0	0 \ 0	90 \ 10	H	0 \ 0	0 \ 0	85 \ 15
M	0 \ 0	50 \ 50	50 \ 5	M	0 \ 0	50 \ 50	50 \ 10	M	0 \ 0	50 \ 50	50 \ 15
L	5 \ 95	5 \ 50	5 \ 5	L	10 \ 90	10 \ 50	10 \ 10	L	15 \ 85	15 \ 50	15 \ 15

P1\P2	H	M	L	P1\P2	H	M	L	P1\P2	H	M	L
H	0 \ 0	0 \ 0	80 \ 20	H	0 \ 0	0 \ 0	75 \ 25	H	0 \ 0	0 \ 0	70 \ 30
M	0 \ 0	50 \ 50	50 \ 20	M	0 \ 0	50 \ 50	50 \ 25	M	0 \ 0	50 \ 50	50 \ 30
L	20 \ 80	20 \ 50	20 \ 20	L	25 \ 75	25 \ 50	25 \ 25	L	30 \ 70	30 \ 50	30 \ 30

P1\P2	H	M	L	P1\P2	H	M	L	P1\P2	H	M	L
H	0 \ 0	0 \ 0	65 \ 35	H	0 \ 0	0 \ 0	60 \ 40	H	0 \ 0	0 \ 0	55 \ 45
M	0 \ 0	50 \ 50	50 \ 35	M	0 \ 0	50 \ 50	50 \ 40	M	0 \ 0	50 \ 50	50 \ 45
L	35 \ 65	35 \ 50	35 \ 35	L	40 \ 60	40 \ 50	40 \ 40	L	45 \ 55	45 \ 50	45 \ 45

Table 1: Possible payoff matrices (demand combinations).

The analysis of the simulations showed that when the differences: H-M or M-L are higher the transition time between the fractious state and the equitable equilibrium is longer. A comparison of the transition time for different payoff matrix is shown in Figure 7.

¹¹ In any case, the sum of the values of L and H is equal to the 100 percent of the pie.




Fig. 7 Number of iterations to equitable equilibrium as a function of L (lowest payoff) and n (number of agents); uncertainty parameter $\varepsilon = 0.1$ and memory length = 10.

3.4 Changing the initial conditions: ‘progressive memory’

In AEY’s model, all the individuals in the experiment had a fixed-size memory along all the matches. Furthermore, their memories contained random values when they are created.

Kubera et al. (2009) explains that it can introduce biases in the results. In our study, we shall suppose that the memory size of each individual grows at a rate of one unity per match, starting with a 0-size memory, until the memory size reaches AEY’s fixed value. The memory size will not grow any longer when it reaches this value.

To fix ideas, let us suppose that we have defined a memory size of 6 ($m=6$). This means that each agent can remember the decision taken by his latest six opponents. In AEY’s original model, a memory size of 6 means that all the agents have six random values in their memories when the game is initialized. After each match, the decision taken by his opponent is stored in his memory, and the oldest value in his memory is eliminated so that the memory size is kept constant.

With our modification of the initial conditions of the system, in the first match, all the agents have six memory positions. However, as they have never played against any other player, their memories are empty. This is the reason why, in the first match, the decision taken by each agent is random. Afterwards, all the agents store the decision taken by their opponents, as they did in AEY’s model. They will use this information to take a decision in the second match, with the same criteria as in AEY’s model. Then, the decision taken by their opponents will be stored in their memories once again. In the third match, each agent will have information about the two previous matches; they will take a decision based on this information and store the decision taken by their opponents, and so on. When the number of matches is higher than the memory size for each agent (m), the agents will store the decisions taken by their opponents in their memories, but will eliminate the oldest value in their memories so that the memory size is equal to m in the following matches.

Figure 8 compares the time it takes for the system to reach the equitable equilibrium, both with and without progressive memory. In the system lacks progressive memory (original AEY's model), agent's memories are initialized with $m=12$ random values. In the case of progressive memory, each agent's memory is started with one random value and their memory grows in one element iteration by iteration until it reaches length $m=12$.




Fig. 8 Comparison of AEY's model with and without progressive memory. Number of iterations to equitable equilibrium. Uncertainty parameter $\varepsilon=0.1$. Memory length = 12.

Although it might seem that the system behaves as Axtell's model, the simulation showed that just by changing the initial conditions, the results of the simulation are completely different.

First, as figure 8 shows, the time it takes for the system to reach the equitable equilibrium is longer than in AEY's original model. Because the first decision is random, the chances of choosing L or H are twice the chances of choosing M, which makes the system approach to the fractious state during the first steps of the simulation. The presence of noise in the system ($\varepsilon \neq 0$), make it possible that agents choose M with certain probability, which leads the system to the equitable equilibrium in the long term. Because of this transitory situation, in which the system tends to approach to the fractious state during some iterations, the number of iterations until the system reaches the equitable equilibrium is higher than in EAY's model.

Secondly, notice that, in the case of progressive memory, the value assigned to ε is crucial. For low values of ε , the system tends to reach a fractious state. The presence of noise makes the agents choose M at some point of the simulation. The increment of the presence of M in their memories makes the agents consider that M is a good reply: eventually, agents learn to compromise and reach an equitable equilibrium. The chances of this situation happening are higher when ε grows.

4 The Model with two agent types (the “tag” model)

In a second experiment, AEY let the agents be distinguishable from one another by introducing a tag: they create two types of agents, each of whom with a different tag (colour). The agents are capable of identifying their opponents’ tag (colour) and they keep the portion of the pie demanded by their opponents in their memories, both with the same and different tag. AEY states that discrimination (segregation) can emerge spontaneously, both when the agents play with other agents of the same type (intra-type matches) and when the agents play against players with different tag (inter-type matches).

To study the different cases of segregation, AEY uses two simplexes: one shows the memory state of the agents when they play against agents with their same tag; the other one display agent memories when they play against agents with different tag. However, our replication did not show any segregation: all the agents learned to compromise and they always reached the equitable equilibrium (independently of their tag)¹².

Then, we tried changing the decision rule, so that the agents chose the best reply against the most frequent option taken by their opponents in previous matches (mode of their memory), see section 3.2. The simulation showed that just after changing the decision rule, segregation emerged spontaneously. In this case, we observed all the possible cases of segregation shown in AEY’s model. We can distinguish two types of segregation: intratype segregation (i.e. discrimination that emerges when players with the same tag play among them) and intertype segregation (i.e. discrimination that arises when agents play against agents with different tag).

Intratyp segregation

Figure 9 shows the three scenarios that can arise when players of the same tag play among them (intratype matches). The percentage of experiments that reach each situation is shown in the table below, for several values of L (lowest payoff).

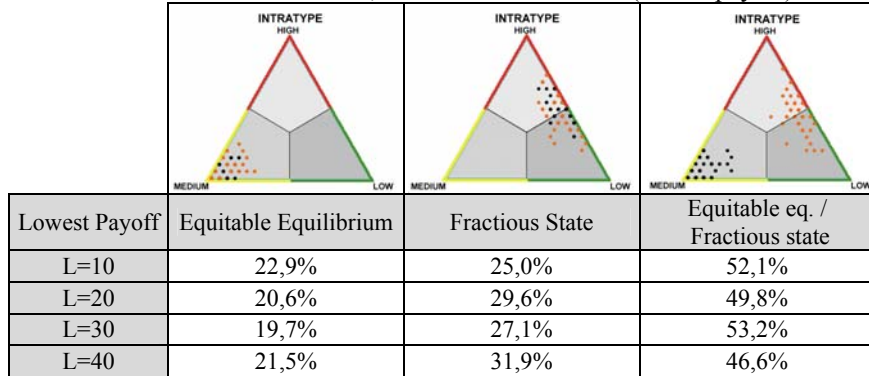


Fig. 9 Intratype memories (memory of the agents when they play against players with the same tag). Modification of AEY’s model with a new decision rule. N=100 (50 agents of each type). M=20. $\epsilon=0$.

¹² We recently asked AEY’s authors about the decision rule that they used to provoke segregation, and we hope their answer would be available for a final version of this work.

In the case of intra-type matches, we could appreciate three different scenarios:

- Equitable equilibrium (all the agents demand M independently of their tag).
- Fractious state (the agents are whether aggressive or passive and do not learn to compromise).
- Intra-type segregation: The agents with one tag reach an equitable equilibrium and the agents with another tag reach a fractious state.

The first and the second scenarios do not show any kind of discrimination: the system reaches an equitable equilibrium or a fractious state independently of the agent's tag, as it did in AEY's model with one agent type. The third scenario is more interesting: when dark players play against dark players, they consider that M is the best response and reach an equitable equilibrium. However, when light players play among them, they do not learn to compromise and the system reaches a fractious state. This happens even though the decision rule is the same for both types of agents.

Intertype segregation

In the case of inter-type matches, we can appreciate the two different scenarios shown in figure 10:

- Equitable equilibrium (all the agents demand M independently of their tag).
- Fractious state (the agents of one colour are aggressive –they choose H– and the agents of another colour are passive –they choose L–).

The percentage of experiments that reach each scenario is listed below.




Fig. 10 Intertype memories (memory of the agents when they play against players with different tag). Modification of AEY's model with a new decision rule. $N=100$ (50 agents of each type). $M=20$. $\varepsilon=0$.

Some of the experiments showed intertype discrimination. When agents of different tags are paired to play, the dark agents find that light agents have frequently demanded H. Consequently, they decide to choose L, which is the only demand that allows them to get a non-zero benefit. On the contrary, after a number of iterations, the light agents have found that light agents are likely to choose low (L). Therefore, they choose (H), as it maximizes their benefit. This situation can be seen as a 'stable fractious state', because the system keeps in this state for long periods of time: all the agents with one tag are aggressive (they all choose H) and all the agents of the other tag are passive (all of them choose L).

The tables below figures 9 and 10 show that, in intratype and intertype games, the chances of reaching one of the scenarios does not change substantially when we modify the payoff matrix.

5 Conclusions

In AEY's model segregation emerges spontaneously as a consequence of the tag recognition. There is not a behaviour rule making agents behave in a different way when they play against agents with their same tag or with different tag. The only difference among the agents is the tag – which a priori does not need to influence on the decision as it is an external property – and the memories about the previous games. Initially, the two types of agents are initialized with the same criteria to get a random memory. After a series of iterations with other agents, they “learn” how to behave depending on whether the agent they meet is same-tag opponent or a different-tag opponent.

The replication of AEY's no-tags model, showed that transition time rises as the memory size and the number of agents grow, as (Axtell et al. 2000) concluded. The simulation of our replication is completely in agreement with their results.

The modification of AEY's no-tags model showed interesting results. We conclude that simple changes within the original model (using the mode instead of the mean to take a decision), provokes dramatic changes in the studied system. In fact, when we introduced this new decision rule, the chances of reaching an equitable equilibrium were considerably lower than in AEY's original model.

Moreover, changing the original payoff matrix in AEY's model resulted in a considerable modification in the transition time: the higher the reward assigned to low, the longer it took for the system to reach the equitable equilibrium.

Initializing the agents with a progressive memory (i.e. a memory which grows one element per match) instead of using AEY's fixed-size memory, showed an interesting scenario: agents tend to be aggressive or passive at first, but after a number of iterations, they learn to compromise. This makes the system reach an equitable equilibrium in the long run. Agents' fractious behaviour in the first stages of the simulation resulted in an increase of the transition time in comparison with AEY's original model.

After replicating the tag model, we conclude that our results are in accordance with the original AEY's work. However, we had to make the assumption that the decision rule used by AEY was to choose the best reply against the most frequent decision taken by the opponents in previous matches (agents use mode and not the mean to endorse the opponent's strategy). If agents consider that opponents will play H/L/M with the same probability that it is recorded in its memory, segregation never emerges in our experiments.

Nevertheless, trying to go a step further, we have inquired if the segregation results are independent of the payoff matrix. We have evaluated different combinations of proportions in the rewards L/H/M and we can conclude that segregation always emerges, independently of the assigned values to these rewards. Furthermore, changing the payoff matrix did not affect the time it took for the system to reach each

scenario (both with and without segregation). This is in contrast with the ‘no-tag model’, in which agents used the mean to maximize their benefit. In this case, it takes much more time to get the equilibrium (both players play M) when the values assigned to L/H/M are not very distant.

In future research we will include new decision rules, such as using moving averages when taking a decision and endorsement mechanisms to assign more relevance to the decisions taken in the recent games than in the older ones. We are currently working in playing the game in a 2D grid and with different social networks topologies, to study how the segregation can affect/be affected when agents are not randomly paired.

Acknowledgments. We acknowledge financial support by the Spanish Council for Science and Innovation: "Agent-based Modelling and Simulation of Complex Social Systems (SiCoSSys)" project, with grant Ref. TIN2008-06464-C03-02.

References

1. Axelrod, R. M. (1986). An Evolutionary Approach to Norms. *American Political Science Review*, 80(4), pp. 1095-1111.
2. Axelrod, R. M. (1997). Advancing the Art of Simulation in the Social Sciences. Published in Rosario Conte, Rainer Hegselmann and Pietro Terna (eds.), *Simulating Social Phenomena*, pp.21-40. Berlin: Springer, 1997.
3. Axtell, R.; Epstein, J.M. and Young, P.(2000). The Emergence of Classes in a Multi-Agent Bargaining Model. *Brookings Institution - Working Papers No.9*. February 2000.
4. Dessalles, J.L., Müller, J.P. and Phan, D. (2007). Emergence in Multi-agent Systems. *Conceptual and Methodological Issues. Agent-Based Modelling and Simulation in the Social and Human Sciences*. 14, pp. 327-355. The Bardwell Press, Oxford.
5. Edmonds, B. and Hales, D. (2003). Replication, replication and replication: Some hard lessons from model alignment. *Journal of Artificial Societies and Social Simulation*, 6(4), 2003.
6. Epstein, J. & Axtell, R. (1996). *Growing artificial societies: social science from the bottom up*. Washington, DC: Brookings and Cambridge, MA: MIT Press.
7. Galan, J. M. & Izquierdo, L. R. (2005). Appearances Can Be Deceiving: Lessons Learned Re-Implementing Axelrod's 'Evolutionary Approach to Norms'. *Journal of Artificial Societies and Social Simulation* 8(3)2. <http://jasss.soc.surrey.ac.uk/8/3/2.htm>
8. Galán J.M., Izquierdo, L.R., Izquierdo, S., Santos J.S., del Olmo, R., López-Paredes, A. and Edmonds, B. (2009). Errors and Artefacts in Agent-Based Modelling. *Journal of Artificial Societies and Social Simulation* vol. 12, no. 1 1. <http://jasss.soc.surrey.ac.uk/12/1/1.html> January 2009.
9. Ito, T.A., Willadsen-Jensen, E.C., & Correll, J. (2007). Social neuroscience and social perception: New perspectives on categorization, prejudice, and stereotyping. In E. Harmon-Jones & P. Winkielman (Eds.), *Social Neuroscience: Integrating Biological and Psychological Explanations of Social Behavior* (pp. 401–421). New York: Guilford.
10. Kubera, Y.; Mathieu, P.; Picault, S. (2009). How to Avoid Biases in Reactive Simulations. 7th International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS'09).

11. López-Paredes A., Hernández C., Pajares J. (2002). Towards a New Experimental Socio-economics. Complex Behaviour in Bargaining. *Journal of Socioeconomics* 31, 423-429, 2002.
12. López-Paredes A., Hernández C., Pajares J. (2004). Social intelligence or tag reasoning?. 2nd ESSA Conference. <http://www.unikoblenz.de/~essa/ESSA2004/files/papers/LopezHernandezPajaresESSA04.pdf>
13. Sansores, C., Pavón, J.: Agent-based simulation replication: A model driven architecture approach. In: Gelbukh, A.F., de Albornoz, A., Terashima-Marín, H. (eds). *MICAI 2005: Advances in Artificial Intelligence*, 4th Mexican International Conference on Artificial Intelligence, Monterrey, Mexico, November 14-18, 2005, Proceedings. Lecture Notes in Computer Science 3789. Springer, Berlin Heidelberg (2005) 244-53.
14. Schelling, T. (1971). "Dynamic Models of Segregation." *Journal of Mathematical Sociology* 1:143-186. 1971.
15. Wilensky, U. and Rand, W. (2007). Making Models Match: Replicating an Agent-Based Model. *Journal of Artificial Societies and Social Simulation*. October 2007.