



Algorithmic Social Sciences Research Unit

ASSRU

Department of Economics
University of Trento
Via Inama 5
381 22 Trento Italy

DISCUSSION PAPER SERIES

8 – 2011/II

AGENT-BASED MODELLING OF THE EL FAROL BAR PROBLEM[♠]

Shu-Heng Chen & Umberto Gostoli

JULY 2011

[♠] Text of the talk given at the ASSRU/Department of Economics Seminar, University of Trento, 31 May, 2011. Professor Shu-Heng Chen is a *Founding Honorary Associate* of ASSRU.

Agent-Based Modeling of the El Farol Bar Problem

Shu-Heng Chen and Umberto Gostoli

AI-ECON Research Center

Department of Economics

National Chengchi University

Taipei, Taiwan

chen.shuheng@gmail.com, u.gostoli@gmail.com

In this paper, we study the self-coordination problem as demonstrated by the well-known El Farol problem (Arthur, 1994), which has later become what is known as the minority game in the econophysics community. While the El Farol problem or the minority game has been studied for almost two decades, existing studies are mostly only concerned with efficiency. The equality issue, however, has been largely neglected. In this paper, we build an agent-based model to study both efficiency and equality and ask whether a decentralized society can ever possibly self-coordinate a result with the highest efficiency while also maintaining the highest degree of equality. Our agent-based model shows the possibility of achieving this social optimum. The two key determinants to make this happen are social preferences and social networks. Hence, not only do institutions (networks) matter, but individual characteristics (preferences) also matter. The latter are open to human-subject experiments for further examination.

Keywords: El Farol Bar problem, Social Preferences, Social Networks, Self-Organization, Emergence of Coordination.

1. Introduction.

The El Farol Bar problem, introduced by Arthur (1994), has over the years become the prototypical model of a system in which agents, competing for scarce resources, deductively adapt their belief-models (or hypotheses) to the aggregate environment that they jointly create. The numerous works that have analyzed and extended this problem along different lines show that perfect coordination, that is, the steady state where the aggregate bar's attendance is always equal to the bar's maximum capacity, is very hard, not to say impossible, to reach, at least under the common knowledge assumption (Fogel, Chellapilla, and Angeline (1999); Edmonds (1999); C. Atilgan, A.

R. Atilgan and Demirel (2008), to name just a few). Works where this assumption has been relaxed, such as those that substituted best-response behavior with reinforcement learning, show that perfect coordination is possible and that it is, indeed, the long-run behavior to which the system asymptotically converges (Whitehead, 2008). However, it is an equilibrium characterized by complete segregation: the population is split into a group of agents who always go (filling the bar up to its capacity all the time) and a group of agents who always stay at home.

In this paper, we pose the question whether a state of perfect coordination with perfect equality, that is, a state where the bar attendance is always equal to its capacity and all the agents go to the bar with the same frequency, can be reached and, if yes, under which conditions. We will refer to this special state as the *socially optimal equilibrium*, as we implicitly assume that, among all states in which the scarce resource is always exploited to the full, the aggregate utility is maximized by its egalitarian division among the agents. In fact, the equality, or fairness, of the outcomes in the El Farol Bar problem is an issue that has been largely neglected by the literature on the subject, with a paper by Farago, Greenwald and Hall (2002) being, to the best of our knowledge, the only exception. However, while this work considers the possibility of reaching a fair outcome through the imposition of a fee by a central planner, in the present paper we consider whether the efficient and fair outcome can emerge from the bottom up, through the process by which the agents' strategies co-evolve and adapt.

Our main finding is that it is possible to reach the socially optimal equilibrium, with the following being two sufficient conditions (although further work is required to assess their necessity). First, the agents need to make use, in their decision-making process, of local information. This means that we have to modify the original model by introducing *social networks*. However, our simulation shows that the social network structure matters: some networks allow the system to reach the socially optimal equilibrium more than others. Second, the agents need to have some kinds of *social preferences*: they need to care about their attendance frequency compared to their neighbors' attendance frequencies.

In the present work we adopt, as a first step, a relatively 'strong' social preferences hypothesis, according to which the agents have the tendency to attend the bar with the same frequency as their neighbors. In fact, the presence of social preferences, in our model, is implicit in the same concept of the socially optimal equilibrium: it is because the agents' utility functions are maximized when their attendance frequencies are equal to each other so that we can define a state where perfect coordination is conjugated with perfect equality as being 'socially optimal'. To better appreciate the effect of social preferences in the outcomes, we also consider

a model with a social network but *without* social preferences. The results show that the presence of social networks is sufficient to allow the system to reach perfect coordination, although in this case equilibria with many different attendance frequency distributions can emerge.

The present paper is organized as follows. In Section 2, we will present a brief review of the literature. In Section 3 we will describe the model and then, in Section 4, we will present the simulations' results. Finally, in Section 5 we will present the conclusions.

2. Previous literature

2.1 The seminal models.

In the El Farol Bar problem, N people decide independently, without collusion or prior communication, whether to go to a bar. Going is enjoyable only if the bar is not crowded, otherwise the agents would prefer to stay home. The bar is crowded if more than B people show up, whereas it is not crowded, and thus enjoyable, if attendees are B or fewer. Arthur assumes that all the agents know the attendance figures in the past m periods and that each of them has a set of k predictors or hypotheses, in the form of functions that map the past t periods' attendance figures into next week's attendance. After each period, the predictors' performance indexes are updated according to the accuracy with which the various predictors forecasted the bar's attendance. Then, the agent selects the most accurate predictor and uses the relative forecast to decide whether to go to the bar or to stay at home during the next period. The characterizing feature of the El Farol Bar problem is that, in such a system, expectations will be forced to differ: if all believe *few* will go, *all* will go, whereas if all believe *most* will go, *nobody* will go, invalidating that belief. Although the competitive process among predictors never comes to rest, it still produces a remarkable statistical regularity: at the macro level, the number of attendees fluctuates around the threshold level B , while, at the micro level, each agent goes B/N percent of the times, in the long run.

Zambrano (2004) shows analytically how the method of inductive inference employed by the agents in Arthur's computer simulation leads the empirical distribution of aggregate attendance to be like those distributions in the set of Nash equilibria of the game. Challet, Marsili and Ottino (2004), in analyzing the El Farol Bar problem with the tools of statistical physics, find that for small m , the relative variance of the fluctuations around the resource level, σ^2/N (that in this, as in other econophysics papers, is taken as a measure of the coordination level), displays a maximum when the agents are endowed with a set of strategies that make them

choose to go to the bar with a frequency equal to the resource level. This suggests that a small bias in the strategies' prescriptions, of either sign, is beneficial as it decreases the fluctuations. They also find that the effect of the distribution of strategies becomes shallower as m increases and it disappears for $m = 6$. However, for large m , the average attendance does not converge to the resource level, so that there is an intermediate memory length which is optimal for the collective behavior that depends on the number of agents N .

Inspired by the El Farol Bar problem, Challet and Zhang (1997) proposed the Minority Game (MG). In the Minority Game there is a population of N (with N being an odd integer) players who have to choose an action (-1 or +1). In each period, the action chosen by the minority wins. The past m outcomes of the game are common knowledge. To choose their next-period state, players use one of their k strategies among a set of strategies drawn at random from the pool of all conceivable strategies, with each strategy being a lookup table assigning an action to any of the past winning actions' configurations. Similar to the El Farol Bar problem, the agents, to make their choice, select their best-performing strategy: after each period, the agents assign a point to all strategies that succeeded in predicting the winning action (regardless of whether they were actually used or not) and zero to the others. The main difference between the two models, apart from the different threshold B (respectively 60 and 50%), is that while in the former no explicit assumption is made regarding the number of agents N (it is set to 100 in Arthur's model), in the MG it is explicitly assumed that N is an odd number, an assumption that, together with the 50% threshold, ensures that there is always a minority side. The simulations show that although the two actions are chosen, as one may expect, 50% of the times in the long run, the average fluctuations' size around the average, a measure of the efficiency with which the system exploits the scarce resource, is inversely proportional to the size of the agents' memory m , at least up to $m \approx 6$. Moreover, simulations show that, contrary to what one may expect, increasing the number of strategies that the players are endowed with, in general, tends to decrease their performance.

In the same paper, Challet and Zhang introduce an 'evolutionary' version of the MG where the worst player is replaced by a new one after some time steps, with the new player being a clone of the best player. To keep some heterogeneity, a mutation process is introduced: one of the best player's strategies is replaced by a new one, randomly drawn from the whole strategies' space. The social learning that takes place in this evolutionary MG makes the average fluctuations' size decrease over time. Moreover, if the memory size m is allowed to change through this evolutionary process, simulations show that an 'arms race' takes place among the players, with the memory size increasing up to a 'saturation' level that increases with the agents'

population size.

However, in spite of the many similarities, the two models differ on one fundamental point: whereas in the MG, as mentioned before, there is always a majority side that makes the wrong choice, in the El Farol Bar problem there is the possibility of exactly hitting the target B , a situation where all the agents, no matter what they decided, made the right choice. This difference makes the *average aggregate payoff* move in opposite directions as we increase the number of agents N : while in the MG the average aggregate payoff is inversely proportional to the aggregate attendance fluctuations' size and, consequently, increases with the number of agents N . In the El Farol Bar problem, it decreases with N . This is because, in this problem, the average aggregate payoff is composed of a negative component, represented by the average aggregate payoff for the times the aggregate attendance is below or above the threshold, and a positive component, represented by the average aggregate payoff for the times the aggregate attendance is exactly equal to the threshold B . Now, while the first component becomes smaller as N increases because of the reduction in the fluctuations' size, the second component also becomes smaller, as with a higher N the probability of hitting the target decreases. The simulations show that the net effect is negative: as we increase N , the positive component becomes smaller more quickly than the negative one, decreasing the average aggregate payoff. Given this different aggregate behavior, in the next sub-section we will focus on the literature on the El Farol Bar problem, considering, within the MG literature, only those papers that have introduced local interaction.

2.2 Extensions: Learning Mechanisms and Local Interactions.

The El Farol Bar problem and the MG have inspired, since their introduction, many works in as many different directions. Here we will focus on two research strands that are relevant to this paper: the introduction of different leaning models in the El Farol Bar problem and the introduction of local interaction in the MG (quite strangely, examples of the adoption of local interaction in the former model and of different learning mechanisms in the latter are much rarer). Among the first research strand, we can distinguish two groups of works: those which retain the best-reply behavior of Arthur's El Farol Bar problem and those which introduce reinforcement learning mechanisms. Among the first Edmonds (1999) proposes an extension of the El Farol Bar problem where agents can change their strategies set by means of a genetic programming (GP) algorithm and are given the chance to communicate with other agents before making their decision of whether to go to the El Farol Bar. Simulations show that although all agents were indistinguishable at the start of the run in terms of

their resources and computational structure, they evolved not only different models but also very distinct strategies and roles.

Another work where the agents' strategies are allowed to co-evolve is that of Fogel, Chellapilla and Angeline (1999). In the model they propose, the agents are endowed with 10 predictors that take the form of autoregressive models with the number of lag terms and the relative coefficients being the variables that evolve over time. For each predictor, one offspring is created (with mutation). The 10 models having the lowest prediction error based on the past 12 weeks of data are selected to be the parent of the next generation. Their simulations show that the system, in a typical trial, has a lower average aggregate attendance (around 56.3%) and a higher standard deviation (17.6) than the ones resulting from Arthur's model.

More recently, Atilgan, Atilgan and Demirel (2008) have explored the effect of (i) the different types of algorithms used by the agents, (ii) the strategy employed to select algorithms from this pool, and (iii) the memory horizon for which attendance data are available to the agents. They show that whether the average attendance will converge to the threshold level or not depends on the algorithm selection procedures of the agents. Changing the algorithm used whenever it fails, irrespective of the past success of the algorithm, and picking up another one randomly, drives the average attendance to the comfort level, as the agents use more information from the past. Taking into account a merit-based stickiness of the algorithms employed in the past exhibits considerable deviation from the path that carries the average attendance to the threshold level. Stickiness not only alters the plateau levels, but also gives rise to large fluctuations on the approach-to-plateau pathways, especially for shorter memory.

Other works have abandoned the best-reply behavior to adopt the more basic reinforcement learning framework. One of the first works where the best-reply behavior of Arthur's original model has been replaced by a kind of reinforcement learning is that of Bell and Sethares (1999). In this paper, the authors present an agent-based model where the agents' strategies are represented by an integer c determining the agents' attendance frequency: if $c = 2$ the agent goes to the bar once every 2 periods; if $c = 3$ he goes once every 3 periods and so on. Every time an agent goes to the bar and has a good time (because the bar was not too crowded) he decreases c (goes more often) whereas, in the opposite case, he increases c (goes less often). No change in the attendance frequency takes place if the agents stay at home, as it is assumed that he cannot assess whether he made the right choice or not.

Subsequently, Franke (2003) proposed a reinforcement learning model that, although quite elaborate, for the purpose of this paper can be summarized as follows: each agent goes to the bar with a probability p . If the bar is not crowded he increases p , while if the bar turns out to be too crowded, he decreases p . If the agent stays at

home, a parameter u determines the extent to which the attendance probability is updated according to the bar's aggregate attendance. In both papers, simulations show that the populations tend to split in two groups: a group of frequent bar-goers and a group of agents who go to the bar very seldom. This result has been analytically obtained by Whitehead (2008). By applying the Erev and Roth (1998) model of reinforcement learning to the El Farol Bar framework, he shows that the long-run behavior converges asymptotically to the set of pure strategy Nash equilibria of the El Farol stage game.

To sum up, the literature on the El Farol Bar problem can be divided in two broad categories, depending on the payoff that, implicitly or explicitly, is assigned to the agents staying at home and this, in turn, is naturally associated with different kinds of learning mechanisms. To make this point clear, let us look at the payoff matrix shown in Figure 1.

| | | |
|---------------|-----------|---------|
| | Uncrowded | Crowded |
| Go to the Bar | G | B |
| Stay at Home | S | H |

Figure 1: El Farol Bar problem payoff matrix

We can distinguish between these two payoff structures:

- a) $G = H > B = S$: this is the typical payoff setting of works adopting the best-reply behavioral model first introduced by Arthur (1994). The payoff, in these models, is represented by the amount by which the strategies' fitness is increased (decreased) after a right (wrong) forecast. To be precise, in some of these models (as the one introduced by Arthur) the strategy fitness is updated by an amount that is inversely proportional to the difference between the aggregate attendance forecasted by the strategy and the actual aggregate attendance. In any case, the characterizing feature of these models is that both the agents going to the bar and those staying at home can assess whether their strategies generated a successful prediction or not (or the extent to which they got close to the target) and can update their strategies' fitness accordingly. In this work, the aggregate attendance fluctuates around an average value that falls between B and a lower bound represented by the average attendance resulting from the agents adopting the mixed strategy that maximizes the aggregate payoff, depending on the values assigned to the many parameters characterizing the best-reply behavior and induction process.

- b) $G > H = S > B$: this is the typical payoff setting of works adopting reinforcement learning. The assumption is that the utility they obtain from staying at home does not depend on the bar attendance. With this payoff structure, two classes of agents emerge: those who often go and those who seldom go to the bar. The learning process will asymptotically lead to a state of perfect coordination with complete segregation, when B agents will always go and N-B agents will always stay at home.

Among the second research strand, based on the introduction of local interaction in the MG, Paczuski, Bassler and Corral (2000) consider a random network of interconnected Boolean elements under mutual influence, the so-called Kauffman network. The performance of each agent is measured by counting the number of times each agent is in the majority. After a certain number of periods, the worst performer, who was in the majority most often, changes his strategy. The Boolean function of that agent is replaced by a new Boolean function chosen at random, and the process is repeated indefinitely. They observe that in some epochs the dynamics of the network takes place on a very long attractor, while, in other epochs, the network is either completely frozen or the dynamics is localized on some attractor with a smaller period. This result, however, appears only when the number of links K is below 6.

Slanina (2000) proposes a model where the agents are placed on a linear chain with nearest-neighbor connections: each agent can 'see' the action and the accumulated wealth of his left-hand neighbor. Each agent is endowed with S strategies. Every agent has a probability p of being an imitator. If an agent is an imitator and his neighbor has larger accumulated wealth than the agent itself, he relegates the decision to the neighbor and takes the same action. In all the other cases (if the neighbor has a lower accumulated wealth or the agent is not an imitator), she will look only at his S strategies and choose the best estimate from them. The results show that there is a local minimum in the dependence of σ^2/N on p , indicating that there is an optimal level of imitation, beyond which the system performs worse. Moreover, this learning dynamics leads to the creation of coherent areas of poor and rich agents.

Finally, Kalinowski, Schulz and Briese (2000) propose a model where the agents are arranged in a circle and everyone obtains the previous decisions of his neighbors as input. For an odd memory m , the decisions of the $(m-1)/2$ left- and right-handed neighbors and the own one are known. The rest of the procedure is the same: each agent looks at the more successful strategy among the s strategies he is endowed with. When all have decided, the minority side is determined, every agent on this side gets a point, the strategies are valued and the next round begins. The simulations show that

the system's efficiency is maximized for $m = 3$. Furthermore, the authors discuss the question of whether the system can be optimized by evolutionary mechanisms. The 'genetic code' of an agent consists of two genes: m and s . After n periods each agent looks at his direct neighbor to the right and to the left and, if the best neighbor has at least 1% more points than the agent, he obtains the properties of this neighbor. Setting $m = 5$ and $s = 4$ as initial states, the simulations show that most agents end up with m and s equal to 2 or 3.

3. The model

For the model we present in this paper, we retain the best-reply strategies of the original El Farol Bar problem. However, we also modify the standard settings by adopting the informational structure introduced by the works on the MG with local interaction. Like in the original El Farol Bar problem, we consider a population composed of 100 agents and set the threshold level to 60 (these values have become commonplace in the literature on this topic). Differing from the original model, each agent can 'see' the actions, the strategies and the strategies' performances of four other agents (his neighbors).

In this paper, we investigate two network typologies (Figure 2): the circular neighborhood, where each agent is connected to the two agents to his left, and the two agents to his right; the von Neumann neighborhood, where the agents occupy a cell in a bi-dimensional grid covering the surface of a torus.

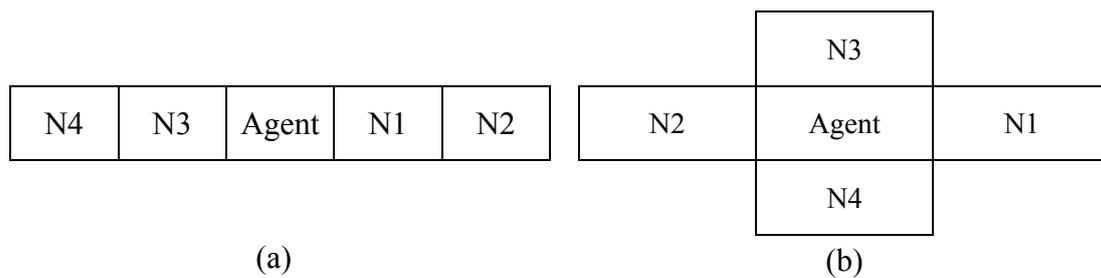


Figure 2: (a) Circular neighborhood; (b) von Neumann neighborhood

Contrary to the prototypical El Farol Bar problem and MG settings, each agent is assigned, at the beginning of the simulation, only one strategy, randomly chosen from the whole strategies' space. This strategy specifies one action to be taken the next period for every possible combination of his four neighbors' actions in the past period. So, the strategies are represented by 16-bit long strings, with a strategies' space of 2^{16} possible strategies (note, at this point, that, with the von Neumann neighborhood, we have the typical settings of *cellular automata*).

The agents are endowed with a memory of length m . Each agent keeps track of his strategy *forecasting success* (that is, the times the strategy prescribed the right action, given the bar's aggregate attendance) and of his *attendance frequency* in the past m periods. The agent's strategy overall performance (or fitness), F is given by (1):

$$F = \frac{\text{forecasting success}}{|\text{attendance frequency} - 0.6| + 1} \quad (1)$$

So, the agent's strategy fitness is greater the higher is its forecasting success (as in the original El Farol Bar problem) and the closer is the agent's attendance frequency to 0.6. Implicit in this rule is the assumption that each agent wants to go to the bar with the same frequency as the other agents. Given this assumption, 0.6 is the only attendance frequency compatible with the bar's capacity: a higher attendance frequency would lead to an overexploitation and a lower frequency would lead to an underexploitation of the scarce resource. The strategy fitness can take any value between 0 and 1.

In any given period, an agent either imitates the strategy of the most successful agent among its neighbors or, with a certain probability p , he mutates his strategy by changing one randomly chosen bit of his strategy. In order for the average attendance associated with any strategy to be computed, it has to be adopted for at least m periods: so, we can think of m as the *trial period* of a strategy. Consequently, an agent changes his strategy (either through imitation or mutation) only if it has been adopted for at least m periods, and, in the imitation process, he considers only those neighbors whose strategy has been adopted for at least m periods. So, to recapitulate, in order for an agent to change his strategy through imitation, five conditions are necessary:

- a) the agent's strategy fitness is below 1.
- b) the agent's strategy is not in its trial period.
- c) The agent has at least one neighbor:
 - whose strategy has a higher fitness than the fitness of the agent's own strategy;
 - whose strategy is not in its trial period;
 - whose strategy is different from the agent's own strategy.

If only the first two conditions are met, the agent, with a probability p , will mutate one rule on his strategy. While the imitation process ensures that the most successful strategies are spread in the population, the mutation process ensures that new, eventually better, strategies are introduced over time. Once the agent has adopted a new strategy (either through imitation or mutation) he will reset his memory to zero and will start keeping track of the new strategy's fitness. Once one agent's strategy

reaches the fitness value of 1, the agent stops the mutation process, as he is perfectly satisfied with his strategy. The socially optimal equilibrium is a state where all the agents' strategies have fitness equal to 1: at this point all the strategies' evolutionary processes stop, as the system has reached the global maximum.

In order to highlight the importance of social preferences for the attainment of the socially optimal equilibrium, we will first show the results of simulations with a model where the social preferences are not present: in this version, the agents just try to develop strategies with good forecasting performances, as in the traditional best-reply framework. In this case, the strategy's fitness F is represented only by the numerator of (1), as the agents have no preferences regarding their attendance frequency.

4. Simulation results

4.1 Social Networks without Social Preferences

In this case, the system, with both social networks, always reaches the perfect coordination, that is, the state where the bar attendance is always equal to the threshold and the agents never make the wrong choice. We observe that the introduction of social networks leads to the emergence of new kinds of equilibrium. We can classify them on the basis of the number of different classes that emerge, each class being characterized by different attendance frequency (the socially optimal equilibrium being the only one with just one class).

Table 1 shows the percentage for each kind of equilibrium (over 1,000 runs), for the circular neighborhood (CN) and the von Neumann neighborhood (vNN). We can see that, with the von Neumann neighborhood, the system has a non-negligible probability of reaching the socially optimal equilibrium (SOE) even if the agents have no social preferences: it is the second most likely equilibrium, with a probability of almost one third of that of the most likely outcome, i.e., 2C.

| | SOE | 2C | 3C | 4C | 5C | > 5C |
|-----|-------|-------|-------|------|-------|------|
| CN | 2% | 62.2% | 15.4% | 3.7% | 13.7% | 2.9% |
| vNN | 18.3% | 53.5% | 15.8% | 3% | 6.3% | 3.1% |

Table 1: Equilibria frequencies without social preferences.

Within the equilibria characterized by the emergence of two classes (2C) the great majority (over 90%) is represented by the well-known 60/40 subdivision between the

agents who always go to the bar and those who always stay at home. The rest (less than 10%) are represented by a new equilibrium characterized by 80 agents with an attendance frequency of 0.5 and 20 agents with an attendance frequency of 1. The vast majority of the equilibria with three classes (3C) are represented by an equilibrium where some agents never go to the bar, some always go and the rest go with an attendance frequency of 0.5. Another relatively frequent outcome is the emergence of five classes (5C). Within this case, the vast majority are represented by a configuration where, besides the three classes mentioned in the 3C case, two groups of agents, that go to the bar with frequencies of 0.4 and 0.6, also emerge.

To sum up, the introduction of social networks leads, on average, to more egalitarian attendance frequency distributions compared to the equilibrium to which the system (asymptotically) converges with reinforcement learning. However, in the absence of preferences regarding the attendance frequency, the less unequal attendance frequency distributions are not associated with higher aggregate utility: the only thing that matters in this setting is the attainment of perfect coordination, and this is reached in every trial.

4.2 Social Networks with Social Preferences

4.2.1 The Circular Neighborhood

Beyond 1,000 simulations, the system always reaches the socially optimal equilibrium. Figure 3 shows the dynamics of the average fitness of the population with the circular neighborhood for a typical run.

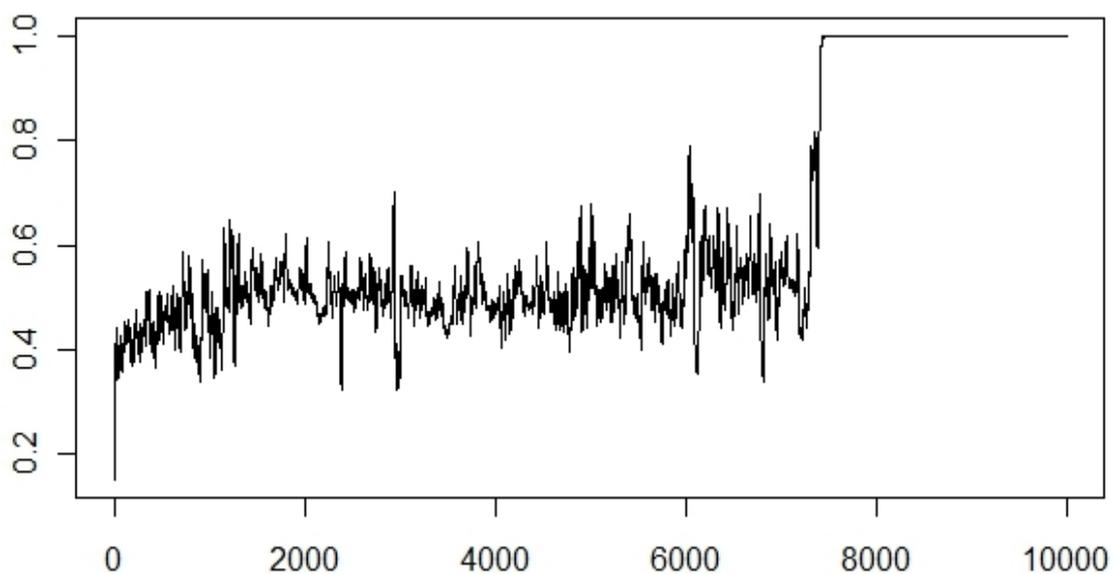


Figure 3: Average fitness dynamics with the circular neighborhood.

No learning process seems to take place: after a period of fluctuations of different sizes around an average value between 0.4 and 0.6, suddenly, around period 7,500, the system jumps to the socially optimal equilibrium.

In Figure 4 we show the distribution of the number of periods the system took to reach the socially optimal equilibrium of over 1,000 trials. We can see that the distribution is highly skewed: while the mode is around 4,000 periods, the average is almost 28,000. However, for a few runs, it took over 150,000 periods to reach the equilibrium.

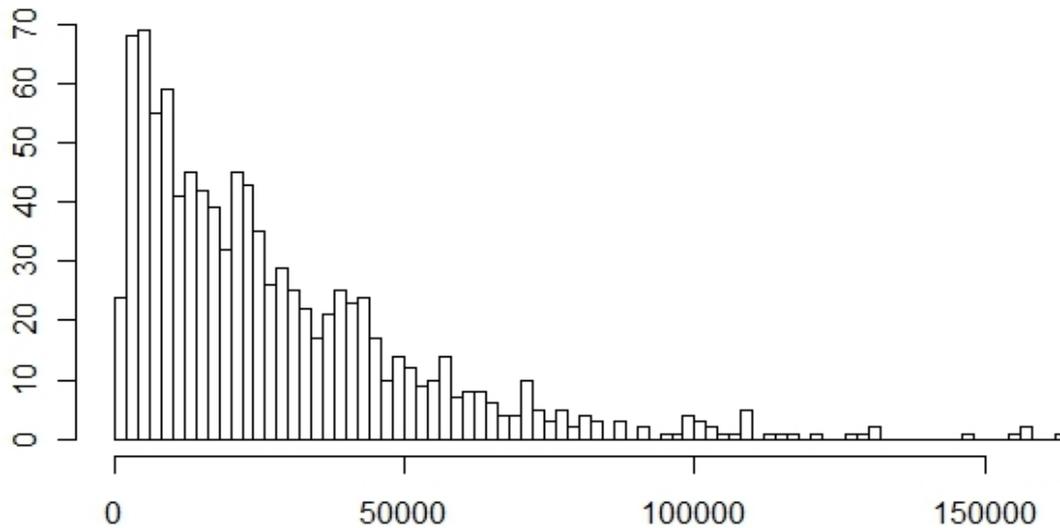


Figure 4: Periods-to-equilibrium distribution with the circular neighborhood

An interesting feature of the model is that, at the equilibrium, the same four strategies always emerge, with both network structures (see Section 4.2.2 for the von Neumann network). Table 2 shows the four strategies emerging with the circular network (in the table, 1 stands for ‘Go’ and 0 stands for ‘Stay at Home’). Even if the whole strategies are composed by 16 binary numbers (that is, 16 rules), at the equilibrium only five combinations repeatedly appear to each agent, so only five of the strategy’s 16 rules are actually used by the agents. From Table 2 we can see that Strategy 1 and Strategy 2, although different, are specified by the same five rules and the same occurs for Strategy 3 and Strategy 4.

| Input | Strategy 1 | Strategy 2 | Input | Strategy 3 | Strategy 4 |
|----------------|------------|------------|-----------------|------------|------------|
| 0-0-1-1 | 0 | 1 | 0-1-0- 1 | 1 | 1 |
| 0-1-1-1 | 0 | 1 | 0-1-1- 0 | 0 | 1 |
| 1-0-1-0 | 1 | 1 | 1-0-0- 1 | 1 | 0 |
| 1-1-0-0 | 1 | 0 | 1-0-1- 1 | 1 | 0 |
| 1-1-0-1 | 1 | 0 | 1-1-1- 0 | 0 | 1 |

Table 2: Emergent strategies with the circular neighborhood

However, the two sets of rules share many similarities as each of the five rules shown in the table on the right side of Table 2 can be obtained from one of the rules shown on the table in the left side of Table 1 by simply moving the first bit of the rule to the last position (for example, in the table on the left, the first rule is 0-0-1-1: moving the first bit in the last position we get 0-1-1-0, the third rule in the table on the right). Moreover, at the equilibrium, the agents do not need to look at all their four neighbors' actions anymore, as each of the four strategies is equivalent to the action of one of the four agents' neighbors: for example Strategy 1 is equivalent to the first neighbor's action (shown in bold in the table on the left); Strategy 2 is equivalent to the third neighbor's action; Strategy 3 is equivalent to the fourth neighbor's action (shown in bold in the table on the left); and Strategy 4 is equivalent to the second neighbor's action. This means that, at the equilibrium, Strategy 1 (if, for example, this is the strategy that emerged), could be substituted by the rule '*Do what the neighbor in position 1 did in the last period*'. Of course, such a simple rule is only able to sustain the cooperation once it has been reached, but can not lead to the equilibrium itself if adopted from a random initial state.

Finally, we observe that all four strategies generate the same 5-period cycle represented by the sequence 1-1-1-0-0. Of course, the strategy the agents are endowed with at the equilibrium makes them follow these cycles asynchronously, so that in every period only 60 of them go to the bar.

4.2.2 The von Neumann neighborhood.

Also in this case, the system reaches the socially optimal equilibrium in all the 1,000 runs performed. Figure 5 shows the dynamics average fitness of the population with the von Neumann neighborhood for a typical run.

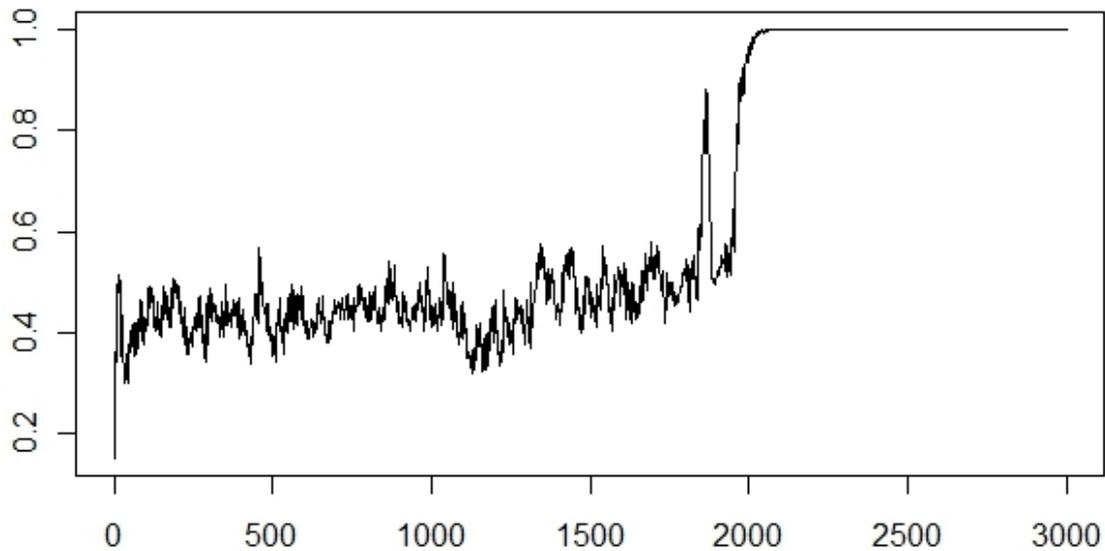


Figure 5: the Average fitness dynamics with the von Neumann neighborhood.

In this chart we can see a spike before the system reaches the perfect coordination. Although not present in every run, these spikes are quite common with the von Neumann neighborhood. They occur when a relatively steady cycle emerges that is characterized by a ‘sea’ of perfectly coordinated agents surrounding an ‘island’ of agents with low fitness. These low-fitness agents are ‘trapped’ in a cycle that causes their attendance frequency to be below or above the optimal value of 0.6. These cycles, however, are short-lived as these agents, sooner or later, will mutate their strategy and some of these mutations will unsettle the cycle in which they had been trapped.

From Figure 6 we can see that, as for the circular neighborhood, the distributions of the periods needed to reach the equilibrium is highly skewed: the mode is around 1,500 periods and the average is around 5,800 periods. In a few runs the system took over 40,000 periods to reach the equilibrium. Overall, compared with the results of the circular neighborhood, the process adopted to reach the equilibrium is between 4 and 5 times faster with the von Neumann neighborhood.

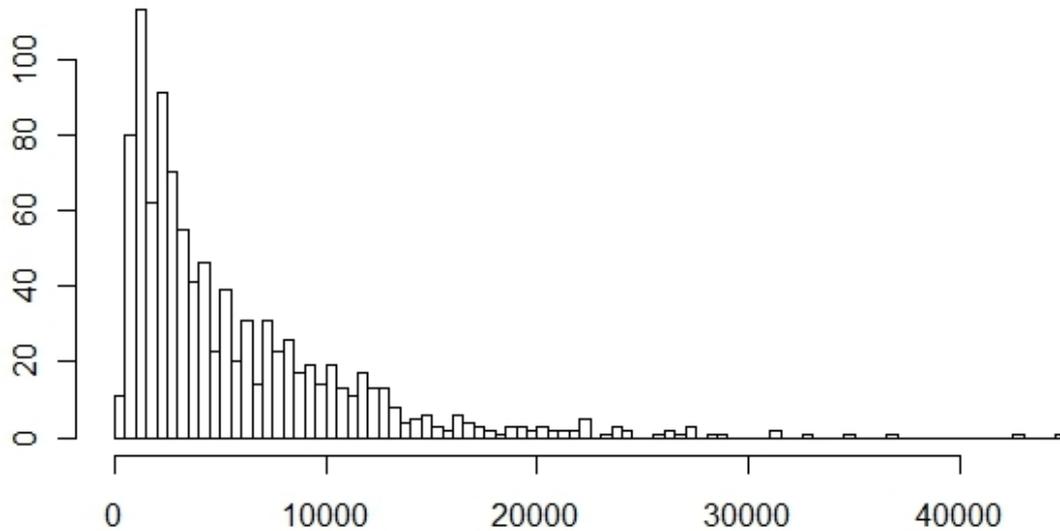


Figure 6: Periods-to-equilibrium distribution with the von Neumann neighborhood

Why is the von Neumann neighborhood a more efficient network structure than the circular neighborhood? A cue to answer this question comes from the observation that one of the crucial parameters for dynamics of the system is the mutation probability p : a few calibration attempts showed that while with the circular neighborhood the average number of periods is minimized with $p \approx 0.002$, with the von Neumann neighborhood it is minimized with $p \approx 0.005$. In other words, the von Neumann neighborhood allows a faster mutation process. We have to keep in mind that in such a system, the mutation process is potentially disruptive: if an optimal strategy is mutated by some agents before it could spread to all the population (because before that point, its fitness is lower than 1), the chances of reaching the equilibrium with this strategy are gone and the system has to wait for another favorable strategy to emerge. So, it is important that every strategy that is introduced in the system spreads quickly, so that the optimality can be checked before the mutation process intervenes to disrupt the coordination ensuing from the spreading strategy. Now, strategies spread more quickly with the von Neumann neighborhood than with the circular neighborhood, given that the average distance between agents is more than 2.5 times lower in the former network compared to the latter (respectively, 5.05 and 12.88). So, the system can effectively explore the highly rugged and ever-changing fitness landscape faster with the former network than with the latter and, consequently, has a greater chance of finding the socially optimal strategy.

Table 3 shows the four strategies emerging with the von Neumann neighborhood. In this case, we can see that, at the equilibrium, only three combinations of actions appear periodically and these are different for Strategy 1 and Strategy 2 on the one hand (the table on the left) and for Strategy 3 and Strategy 4 on the other (the table on

the right). Also in this case, at the equilibrium, the agents do not need to look at all their neighbors' actions to follow their strategy.

| Input | Strategy 1 | Strategy 2 | Input | Strategy 3 | Strategy 4 |
|----------------|------------|------------|----------------|------------|------------|
| 0-1-0-1 | 0 | 1 | 0-1-1-0 | 0 | 1 |
| 1-0-1-0 | 1 | 0 | 1-0-0-1 | 1 | 0 |
| 1-1-1-1 | 1 | 1 | 1-1-1-1 | 1 | 1 |

Table 3: Emergent strategies with the von Neumann neighborhood

While in the case of the circular neighborhood, every strategy was associated with a neighbor, in this case every strategy is associated with *two* neighbors that the agent can equivalently look at: Strategy 1 is equivalent to the actions of the first and the third neighbor (shown in bold in the table on the left); Strategy 2 is equivalent to the actions of the second and the fourth neighbor; Strategy 3 is equivalent to the actions of the first and the fourth neighbor (shown in bold in the table on the right); and Strategy 4 is equivalent to the actions of the second and the third neighbor. So, at the equilibrium, if, for example, Strategy 1 emerges, this strategy can be substituted by the rule ‘*Do what your neighbor in position 1 or your neighbor in position 3 did last period*’. If we look at the von Neumann neighborhood in Figure 1, we can notice that the two neighbors associated with each strategy are adjacent to each other.

Finally, we observe that, while all the four strategies that emerge with the circular neighborhood generate the same 5-period cycle 1-1-1-0-0, with the von Neumann neighborhood, besides this 5-period circle, two 10-period cycles emerge: the cycle 1-1-1-1-0-0-1-1-0-0 and the cycle 1-1-0-0-1-1-0-1-1-0. The three cycles can be generated by any of the four strategies that emerge at the equilibrium.

5. Conclusions

While the El Farol Bar problem has been studied for almost two decades, existing studies are mostly concerned with efficiency only. The equality issue, however, has been largely neglected. In this paper, we present an agent-based model to assess whether a decentralized society can ever possibly self-coordinate a result with the highest efficiency while also maintaining the highest degree of equality. The main differences between the model we presented in this work and those previously introduced are two: a) the presence of social networks through which the agents can access the information regarding their neighbors' choices; and b) the introduction of social preferences, in the form of a fitness function that takes into account the frequency with which the agents go to the bar. The reason behind the introduction of

social networks is the idea that only by using local information do the agents have the chance to coordinate: if they use the aggregate information, as in the best-reply models proposed so far, too many or too few of them are likely to go to the bar, as the aggregate information generates herding behavior. As regards the introduction of social preferences, their presence is implicit in our definition of a socially optimal outcome.

The simulations show that, with these assumptions, the system always reaches a state where all the agents go to the bar 60% of the time and the bar is always filled to its capacity although it could take many periods to reach such a state. The simulations also show that, at the equilibrium, the strategies generate 5-period or 10-period cycles that all the agents follow in an asynchronous way so that their aggregate attendance is always equal to the threshold.

Moreover, the simulations show that there are some kinds of network structures that process the information more efficiently, allowing the system to reach the socially optimal equilibrium much faster than others. In particular, the von Neumann neighborhood seems more likely to lead the system to the socially optimal outcome than the circular neighborhood, and this is true even if we remove social preferences (although in this case this outcome could not be properly defined as ‘socially optimal’ anymore). Simulations with the model without social preferences show that the presence of social networks allows the system to always reach a state of perfect coordination, with many different kinds of different attendance frequency distributions.

The extent to which social networks can enhance coordination among real subjects and whether their behavior is affected by social preferences are issues open to further examination.

Acknowledgements

NSC research grant no. 98-2410-H-004-045-MY3 is gratefully acknowledged.

References

- Arthur, B. W. (1994) . Complexity in economic theory. Inductive reasoning and bounded rationality. *American Economic Review*, 82(2), 406-411.
- Atilgan, C., A. R. Atilgan and G. Demirel (2008). Collective behavior of El Farol attendees. *Advances in Complex Systems*, 11(4), 629–639.
- Bell, A. M., and W. A. Sethares (1999). The El Farol problem and the internet: Congestion and coordination failure. Society for Computational Economics: Computing

in Economics and Finance Working Paper, No. 812.

- Challet, D., and Y.-C. Zhang (1997). Emergence of cooperation and organization in an evolutionary game. *Physica-A* 246 (3-4), 407-18.
- Challet, D., M. Marsili, and G. Ottino (2004). Shedding light on El Farol. *Physica-A: Statistical Mechanics and its Applications*, 332(1), 469-482.
- Edmonds, B. (1999). Modelling bounded rationality in agent-based simulations using the evolution of mental models, in *Computational Techniques for Modelling Learning in Economics*, T. Brenner, ed. Dordrecht: Kluwer, 305–332.
- Erev, I., and A. E. Roth (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4), 848-881.
- Farago, J., A. Greenwald, and K. Hall (2002). Fair and efficient Solutions to the Santa Fe Bar Problem. Proceedings of the Grace Hopper Celebration of Women in Computing.
- Fogel, D. B., K. Chellapilla and P. J. Angeline (1999). Inductive reasoning and bounded rationality reconsidered. *IEEE Transactions on Evolutionary Computation* 3(2): 142-146
- Franke, R. (2003). Reinforcement Learning in the El Farol Model. *Journal of Economic Behavior and Organization*, 51, 367-388.
- Kalinowski, T., H.-J. Schuklz and M. Briese (2000). Cooperation in the minority game with local information. *Physica-A*, 277, 502-508.
- Paczuski, M., K. E. Bassler and A. Corral (2000). Self-organized networks of competing Boolean agents. *Physical Review Letters*, 84(14), 3185–3188.
- Slanina, F. (2000). Social organization in the minority game model. *Physica-A*, 286, 367-376.
- Whitehead, D. (2008). The El Farol Bar problem revisited: Reinforcement learning in a potential game. Edinburgh School of Economics, University of Edinburgh, ESE Discussion Papers, 186.
- Zambrano, E. (2004). The interplay between analytics and computation in the study of congestion externalities: The case of the El Farol problem. *Journal of Public Economic Theory*, 6(2), 375-395.