

When the El Farol Bar Problem Is Not a Problem? The Role of Social Preferences and Social Networks

Shu-Heng Chen and Umberto Gostoli *

AI-ECON Research Center - National ChengChi University

Abstract. In this paper, we study the self-coordination problem as demonstrated by the well-known El Farol Bar problem (Arthur, 1994), which later becomes what is known as the Minority Game in the econo-physics community. While the El Farol problem or the minority game has been studied for almost two decades, existing studies are most concerned with efficiency only. 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 highest efficiency while also maintaining a 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 does institution (network) matters, but individual characteristics (preferences) also matters. The latter part is open for human-subject experiments for further examination.

1 Introduction

The El Farol Bar problem, introduced by Arthur (1994) has become over the years the prototypical model of a system in which agents, competing for scarce resources, adapt inductively their belief-models (or hypotheses) to the aggregate environment they jointly create. The numerous works that have analyzed and extended along different lines this problem 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 Angelina (1999); Edmonds (1999), 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 split into a group of agents who always go (filling the bar up to its capacity all the times) and a group of agents who always stay at home.

* Corresponding author. E-mail: u.gostoli@gmail.com

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 exploited to the full, the aggregate utility is maximized by its equalitarian division among all 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 of Farago, Greenwald and Hall (2002) being, to the best of our knowledge, the only exception. However, while in this latter work the authors consider the possibility to reach 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. Second, the agents need to have some kinds of social preferences: they need to care about their attendance frequency compared with 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 of their neighbors. Moreover, we find that the number of periods it takes for the system to reach the socially optimal equilibrium depends on the kind of social network in which the agents are embedded: with some networks structures, the process leading to the equilibrium is faster than with others.

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

2 Previous Literature

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 at 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. All the agents know the attendance figures in the past m periods and each of them has a set of k predictors, in the form of functions that map the past m periods' attendance figures into next week's attendance. After each period, the predictors' performance indexes are updated according to the accuracy with which the predictors forecasted the bar's attendance. Then, the agent selects the most accurate predictor and uses the relative

forecast to decide whether going to the bar or staying at home in the next period.

Inspired by the El Farol Bar problem, Challet and Zhang (1997) proposed the Minority Game (MG). The main difference among the two models, apart from the different ratio B/N (respectively 0.6 and 0.5), is that while in the former no explicit assumption is made regarding N (in the Arthur's model it is set to 100), 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 majority of agents making the wrong choice. The characterizing feature of these two models is that, in every period, a large share of agents in every period will make a wrong forecast and, consequently, the set of strategies actually adopted by the population will keep changing all the times. Although the competitive process among predictors never comes to rest, the system's dynamics is characterized by a remarkable statistical regularity: at the macro level, the number of attendees fluctuates around the threshold level B/N . On the other hand, and most importantly in his paper's context, these models show that, in systems with heterogeneous sets of expectations and information sets that are common knowledge, such as those described above, it is impossible (at least from a statistical point of view, in the case of El Farol Bar problem) to reach a perfectly coordinated state with no fluctuations.

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 for 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 the Arthur's El Farol Bar problem and those adopting some kind of reinforcement learning mechanism. In the first group, Edmonds (1999) proposes an extension of the El Farol Bar problem where agents can change their strategies set by the means of a genetic programming (GP) algorithm and are given the chance to communicate with other agents before making their decision whether to go to the bar. Simulations show that, although all agents were indistinguishable at the start 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 in the past 12 periods 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 the Arthur's model.

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 the original Arthur's 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 elaborated, 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 these last two 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.

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 with 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, otherwise, the network is either completely frozen or the dynamics is localized on some attractor with a smaller period. 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 her 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 her neighbor has larger accumulated wealth than the agent itself, she 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 her 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 perform

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 on a circle and everyone gets the previous decisions of his neighbors as input. The decisions of the $(m - 1)/2$ left and right handed neighbors and the own one are known. 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. Simulations show that the system's efficiency is maximized for $m = 3$. Furthermore, the authors optimize the system through an evolutionary mechanism. 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 gets the properties of this neighbor. Simulations show that, setting $m = 5$ and $s = 4$ as initial states, most agents end up with m and s equal to 2 or 3.

3 The Model

In the model we present, we retain the best-reply strategies of the original El Farol Bar problem, however we 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 by $N = 100$ agents and set the attendance threshold $B/N = 0.6$. Each agent can 'see' the actions, the strategies and the strategies' performances of four other agents (his neighbors, indicated with N1, N2, N3 and N4). In this paper, we investigate two network typologies shown in Fig. 1: 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*, with the agents occupying a cell in a bi-dimensional grid covering the surface of a torus.

Contrarily to the prototypical El Farol Bar problem and MG settings, each




Fig. 1. Circular (a) and von Neumann (b) neighborhoods

agent is assigned, at the beginning of the simulation, only one strategy (that is, $k = 1$), randomly chosen from the whole strategies' space. The strategy specifies

the action the agent has to take in the current period for every possible combination of his four neighbors' actions in the previous 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*). We define the variable $d_i(t)$ as the action taken by agent i in period t : it takes the value 1 if the agent goes to the bar and the value 0 otherwise. Moreover, we define the variable $s_i(t)$ as the outcome of agent i 's decision in period t : it takes value 1 if the agent took the *right* decision (that is, if he went to the bar and the bar was not crowded or if he stayed at home and the bar was too crowded) and it takes value 0 if the agent took the *wrong* decision (that is, if he went to the bar and the bar was too crowded or if he stayed at home and the bar was not crowded). The agents are endowed with a memory of length m . This means that they store in two vectors, \mathbf{d} and \mathbf{s} of length m , the last m values of d and s , respectively. So, at the end of any given period t , agent i 's vectors \mathbf{d}_i and \mathbf{s}_i , are composed, respectively, by $d_i(t), d_i(t-1), \dots, d_i(t+1-m)$, and by $s_i(t), s_i(t-1), \dots, s_i(t+1-m)$. Agent i 's *attendance frequency*, a_i , is defined by (1):

$$a_i = \frac{1}{m} \sum_{j=t}^{t+1-m} d_i(j) \quad (1)$$

The attendance frequency's value can go from 1, if the agent always went to the bar, to 0, if the agent never went to the bar, in the last m periods. Moreover, the agent i 's *forecasting success*, f_i , is given by (2):

$$f_i = \frac{1}{m} \sum_{j=t}^{t+1-m} s_i(j) \quad (2)$$

The forecasting success's value can go from 1, if the agent always made the right choice, to 0, if the agent always made the wrong choice, in the last m periods. Finally, we define the *strategy's fitness* of agent i , F_i , as follows:

$$F_i = \frac{f_i}{|a_i - 0.6| + 1} \quad (3)$$

We can see that the strategy fitness's value can go from the minimum value of 0, if the agent's *forecasting success* is 0, to the maximum value of 1, if the agent's *forecasting success* is 1 and the agent's *attendance frequency* is 0.6. Implicit in this rule is the assumption that each agent wants to go to the bar with the same frequency of the other agents. Given this assumption, an attendance frequency of 0.6 is the only one compatible with the full exploitation of the bar's capacity: a higher attendance frequency would lead to an over-exploitation while a lower attendance frequency would lead to the under-exploitation of the bar's capacity.

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 any strategy to take part to the selection and replication process, it has to be

adopted for at least m periods: so, we can think of m as the *trial* period of a strategy. This means that an agent changes its 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 its 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 the first two conditions are met but at least one of the other three is not (that is, if the agent has not yet reached the optimal strategy but in the current period he cannot imitate any of his neighbors), the agent, with a probability p , will mutate one rule of his strategy. While the imitation process ensures that the most successful strategies 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 imitation and mutation processes, as it is perfectly satisfied by its 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.

4 Simulations' Results

Fig. 2 shows the dynamics of the *average fitness* of the population in a typical run for each of the two social networks considered: the von Neumann neighborhood (vNN) and the circular neighborhood (CN). As we can see, the system reaches the socially optimal equilibrium (where the average fitness equal to 1) in both cases, but with the von Neumann neighborhood the process appear to be faster than with the Circular neighborhood. From Fig. 2, looking at the dynamics of the average fitness with the von Neumann neighborhood, we can see a spike just before the system reaches the perfect coordination, with the average fitness jumping above 0.8 before falling back to a level close to 0.5. Although not present in every run, these spikes are quite common with the von Neumann neighborhood. They occur when a relatively steady cycle emerges characterized by a 'sea' of perfectly coordinated agents surrounding an 'island' of low-fitness agents with attendance frequencies below or above the optimal value of 0.6. These cycles are relatively steady because the agents at the border of the 'island' have the same strategy of the perfectly coordinated agents, so cannot change their strategy by

imitation. However, sooner or later, their strategy will change by mutation and some of these mutation will unsettle the cycle in which they had been trapped.

Fig. 3 shows the distribution of the periods the system takes to reach the




Fig. 2. The average fitness dynamics

socially optimal equilibrium with the two social networks (note that the scales of the x-axes are different for the two cases). First of all, we observe that over 1000 simulations we run for each kind of network, the system always reached the socially optimal equilibrium. Second, it takes on average around 5 times less to reach the socially optimal equilibrium with the von Neumann neighborhood than with the circular neighborhood (around 5800 in the former case periods against 28500 periods in the latter).

An interesting feature of the model is that, at the equilibrium, always the




Fig. 3. Circular (a) and von Neumann (b) neighborhoods

same four strategies emerge, with both network structures. Fig. 4 shows the four strategies emerging with the circular network. Even if the whole strategies are

composed by 16 binary numbers, corresponding to the 16 possible combinations of the agent's four neighbors' actions, at the equilibrium a cyclical pattern composed by five combinations emerges, so only five of the strategy's 16 rules are used by the agents. Moreover, looking at Fig. 4 we can see how, at the equilibrium, the agents do not need to look at all their four neighbors' actions anymore, as each of the four emerging strategies is equivalent to the action of one of his neighbors: Strategy 1 is equivalent to the N1's action (shown in bold); Strategy 2 is equivalent to the N3's action; Strategy 3 is equivalent to the N4's action (shown in bold; Strategy 4 is equivalent to the N2's action. In other words, if, for example Strategy 1 emerges, the socially optimal equilibrium is maintained with the agents following the rule "Do what your neighbor N1 did in the last period", and similarly for the other three strategies.

| Input | Strategy 1 | Strategy 2 |
|-----------------|------------|------------|
| 0 -0-1-1 | 0 | 1 |
| 0 -1-1-1 | 0 | 1 |
| 1 -0-1-0 | 1 | 1 |
| 1 -1-0-0 | 1 | 0 |
| 1 -1-0-1 | 1 | 0 |

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

Fig. 4. Emergent strategies with the circular neighborhood

Fig. 5 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. Also in this case, at the equilibrium, the agents do not need to look at all their neighbors' actions to follow their strategy. 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. If we look at the von Neumann neighborhood of Fig. 1, we can see that each of the four strategies corresponds to the actions of two adjacent neighbors: the neighbor N1 or N3 for Strategy 1 (shown in bold); the neighbor N2 or N4 for Strategy 2; the neighbor N1 or N4 for Strategy 3 (shown in bold); the neighbor N2 or N3 for Strategy 4. So, like with the circular neighborhood case, also in this case all the agents need to do in order to maintain the socially optimal equilibrium, once it has been reached, it is to follow simple rules based on the previous action of just one of their neighbors: if, for example, Strategy 1 emerges, the rule to follow is "Do what your neighbors N1 or N3 did in the last period", and similarly for the other three strategies. We have to note that there is no guarantee, and indeed it is very unlikely, that the system would have ever reached the socially optimal equilibrium if the agents were to follow these simple rules from the beginning.

Finally, we observe that all the four strategies that emerge with the circular neighborhood, generate the same 5-period cycle represented by the cycle

| Input | Strategy 1 | Strategy 2 |
|----------------|------------|------------|
| 0-1-0-1 | 0 | 1 |
| 1-0-1-0 | 1 | 0 |
| 1-1-1-1 | 1 | 1 |

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

Fig. 5. Emergent strategies with the circular neighborhood

[1-1-1-0-0], whereas with the von Neumann neighborhood, beside 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].

5 Conclusions

While the El Farol Bar problem has been studied for almost two decades, existing studies are most 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 highest efficiency while also maintaining a highest degree of equality, a steady state that we called *socially optimal equilibrium*. Our agent-based model shows the possibility of achieving this equilibrium with the following being two sufficient conditions: a) the agents have to take their decisions on the basis of local information (i.e. their neighbors' past attendances): global information does not allow to reach the socially optimal equilibrium because it causes herd behavior, causing too many or too few people going to the bar at the same time; b) the agents need to have a preference for an equal attendance: if the agents are indifferent to whether they go less or more than their neighbors, the system is likely to converge to an equilibrium where 60% of agents always go and 40% never go to the bar.

Moreover, the simulations results showed that once the socially optimal equilibrium has been reached, it can be maintained through the adoption of rule-of-thumb strategy allowing the agents to minimize their decisional workload. This fact warns us of the difficulty to infer the evolutionary processes that led to the emergence of behavioral rules, just as would have been very difficult, if not impossible, for an observer, looking at the simple behavioral rules emerging after the socially optimal equilibrium has been reached, to infer the complex co-evolutionary process that led to the equilibrium itself. The observation of those simple rules, in fact, would even mislead him in his search for the process leading to the equilibrium, as no equilibrium would have ever been reached if the agents were to follow these rules from the beginning.

Finally, our simulations show that the network structure connecting the agents is represented by the von Neumann neighborhood, the system reaches the socially optimal equilibrium five times faster (on average) than if the agents are connected through a circular network: it seems that different social networks

process the information with different levels of efficiency. So we can say that both the individual preferences and institutional framework matter as regards to the kind of social convention that finally emerges and the time it takes for it to emerge.

References

- Arthur, B.W. (1994). Complexity in economic theory: Inductive reasoning and bounded rationality. *American Economic Review*, 82(2), 406–411.
- Bell, A.M., & Sethares W.A. (1999). The El Farol problem and the Internet: Congestion and coordination failure. *Computing in Economics and Finance Working Paper*, 1999 812. Society for Computational Economics.
- Challet, D., & Zhang, Y.-C. (1997). Emergence of cooperation and organization in an evolutionary game. *Physica-A*, 246(3-4), 407-18.
- Edmonds, B. (1999). Gossip, sexual recombination and the El Farol bar: Modelling the emergence of heterogeneity. *Journal of Artificial Societies and Social Simulation*, 2(3).
- Erev, I., & Roth, A.E. (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., Greenwald, A., & Hall, K. (2002). Fair and efficient solutions to the Santa Fe Bar problem. *Proceedings of the Grace Hopper Celebration of Women in Computing*.
- Fogel, D. B., Chellapilla, K., Angeline, P.J. (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(3), 367–388.
- Kalinowski, T., Schuklz, H.-J., & Briese, M. (2000). Cooperation in the minority game with local information. *Physica-A*, 277(3-4), 502–508.
- Paczuski, M., Bassler, K.E., & Corral, A. (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(1-2), 367–376.
- Whitehead, D. (2008). The El Farol Bar problem revisited: Reinforcement learning in a potential game. *ESE Discussion Papers*, 186, Edinburgh School of Economics, University of Edinburgh.