Laboratory of Economics and Management
Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà, 33 - I-56127 PISA (Italy)
Tel. +39-050-883-341 Fax +39-050-883-344
Email: lem@sssup.it Web Page: http://www.lem.sssup.it

LEM
Working Paper Series

Minority Games, Local Interactions, and Endogenous Networks

Giorgio Fagiolo[*]

and

Marco Valente[+]

[*] *Sant'Anna School of Advanced Studies, Pisa, Italy*
+*University of L'Aquila, Faculty of Economics, L'Aquila, Italy*

2004/17 **September 2004**

ISSN (online) 2284-0400

Minority Games, Local Interactions, and Endogenous Networks*

Giorgio Fagiolo† Marco Valente‡

January 10, 2005

Abstract

We study a local version of the Minority Game, where agents are placed on the nodes of a directed graph. Agents care about being in the minority of the group of agents they are currently linked to and employ myopic best-reply rules to choose their next-period state. We show that, in this benchmark case, the smaller the size of local networks, the larger long-run population-average payoffs. We then explore the collective behavior of the system when agents can: (i) assign weights to each link they hold and modify them over time in response to payoff signals; (ii) delete badly-performing links (i.e., opponents) and replace them with randomly chosen ones. Simulations suggest that, when agents are allowed to weigh links but cannot delete/replace them, the system self-organizes into networked clusters that attain very high payoff values. These clustered configurations are not stable and can easily be disrupted, generating huge subsequent payoff drops. If, however, agents can (and are sufficiently willing to) discard badly performing connections, the system quickly converges to stable states where all agents get the highest payoff, independently of the size of the networks initially in place.

Keywords: Minority Games, Local Interactions, Endogenous Networks, Adaptive Agents.

JEL Classification: C72, C73.

1 Introduction

In the last years, both physicists and economists have become increasingly interested in investigating the collective properties of dispersed dynamical systems composed of many

*Thanks to David Belsley, Ulrich Horst, Alan Kirman, David Lane, Mario Lavezzi, Luigi Marengo, Matteo Marsili, Gérard Weisbuch, and an anonymous referee, for their helpful and stimulating comments.

†Corresponding Author. Faculty of Economics, University of Verona, Verona (Italy) and Sant'Anna School of Advanced Studies, Pisa (Italy). Address: Piazza Martiri della Libertà, I-56127 Pisa (Italy). Tel: +39-050-883341. Fax: +39-050-883344. Email: fagiolo@sssup.it

‡Faculty of Economics, University of L'Aquila, L'Aquila (Italy). Email: mv@business.auc.dk

boundedly-rational agents who directly interact over time (Kirman, 1997).

A well-known instance of such a system is the Minority Game (MG), first introduced as a model of inductive rationality in the famous "El-Farol Bar Problem" by W. Brian Arthur (Arthur, 1994) and then explored in details by Challet and Zhang (1997).

In a nutshell, the standard formulation of the MG envisages a population of N (odd) players who repeatedly choose a binary state (-1 or $+1$). In each time period $t = 1, 2, ...$, the state chosen by the minority wins. Agents who are in the minority get a point, the others get zero. Agents are only allowed to observe the last $m \geq 1$ winning sides: history is the only common information. To choose their next-period state, players use one of their $k \geq 2$ strategies, each strategy being a lookup table assigning an output (i.e., the state to be chosen in the next period) to any of the 2^m possible state configurations. Agents always select their best-performing strategy to make their choice. The performance of any given strategy evolves through time. Agents are initially endowed with a random *repertoire* of strategies, drawn at random from the pool of all 2^{2^m} conceivable ones. If an agent is successful in a given time period, it assigns a point to all strategies that would have predicted the winning state – no matter if they were actually used or not – and zero otherwise.

The standard MG has become a paradigm for studying systems where adaptive agents compete for scarce resources and has been recently employed to study the dynamics of stock markets and market-entry games (see, *inter alia*, Challet, Marsili, and Zhang (2000), Marsili and Challet (2001), Bottazzi, Devetag, and Dosi (2003) and Ochs (1990, 1995)). From a theoretical perspective, the standard MG model has been extensively studied both numerically and analytically (Challet and Zhang, 1998) and a huge number of contributions have been exploring a large spectrum of possible extensions (see the Minority Game's web page http://www.unifr.ch/econophysics/minority in the Econophysics Forum Internet Site for an exhaustive annotated bibliography).

In particular, a recent stream of research has investigated the consequences of relaxing two key assumptions of the basic framework, namely *global interactions* (each agent cares about being in the minority of the whole population) and *common information* (all players have access to the same information – the globally winning side). The underlying idea is that each agent playing the MG could instead have access only to a local source of information, for example the state played in the past by those agents who are the "closest ones" in some underlying socio-economic space (Fagiolo, 1998). For example, Paczuski, Bassler, and Corral (2000) models a MG where agents are placed in a Kauffman NK-network (Kauffman, 1993). Agents can only observe the state of the individuals they are currently linked with and hold only one strategy. The latter maps the past state of one's neighbors into the state to be chosen in the next period. Similarly, Kalinowski, Schulz, and Briese (2000) and Slanina (2000) place their agents over one-dimensional boundaryless lattices (circles) and allow them to observe the states held in the recent past by their nearest-neighbors. Notice, however, that in these models agents interact *locally* but play *globally* the MG: players are indeed rewarded only if they choose the *globally* winning side. Conversely, Moelbert and De Los Rios (2002) studies a local MG in which the population is spatially distributed over a regular lattice with nearest-neighbor interactions. Agents only observe time $t-1$ states played by their neighbors and care about being in the minority of their *local* group.

Although the study of "local" versions of the standard setup has been shedding some light on the ways in which heterogeneous information affects collective MG dynamics, very little attention has been paid to the role played by network structures and local interactions in shaping long-run aggregate outcomes. For example, the existing literature has extensively addressed the study of MGs played over very regular network structures such as homogeneous lattices. On the contrary, the consequences of assuming more asymmetric structures – such as generic, random directed graphs – are still poorly investigated (see Kirley (2004) for an exception in the case of scale-free networks).

Furthermore, local MGs have almost exclusively focused on "frozen" or "static" interaction structures. This means that agents always interact with the same opponents and do not have the faculty to modify the structure of their network *endogenously* (see, however, Anghel, Toroczkai, Bassler, and Korniss (2004) and Lo, Chan, Hui, and Johnson (2004) for two examples of exogenous network dynamics). Nevertheless, endogenous network formation is well-known to crucially alter the properties of collective dynamics in spatial games (see for example the discussion in Fagiolo, Marengo, and Valente (2004) and Fagiolo (2005)).

In this paper, we begin to explore in more detail the role played by networks (and the evolution thereof) in local MGs.

We consider a population of N agents living in a discrete-time economy. Each agent is initially connected through unilateral links (directed edges) to a fraction of other players, whose last-period states are the only information they are allowed to observe. Agents are rewarded with one point if they are successful – that is, if they play as the (strict) minority of their local group does – and zero otherwise.

In any time period, only a fraction of the agents are allowed to consider whether to change their state. Agents decide their next-period state by evaluating the state of their local network.

We implement two network evaluation setups. In the first one, we assume *non-weighted links*: each opponent always counts as one, no matter what it did in the past. Therefore, agents simply choose the state played by the strict minority of their peers. In the second setup, we introduce *weighted links*. Each agent separately evaluates every link (opponent) and assigns to it a weight that increases only if the linked agent suggested the minority side in the past. State decisions will then depend upon a comparison between the sum of weights associated to agents playing $+1$ and -1 in the local group. We call this second setup a "weighted" minority game (WMG).

As far as network dynamics is concerned, we simulate the behavior of the system in two scenarios. The first scenario assumes *frozen networks*. The network initially in place cannot be modified, and only their weights might possibly evolve. In the second scenario *endogenous networks* are assumed: agents can decide whether to remove a badly performing link, that is a link whose weight becomes smaller than a given threshold. Since agents do not have good knowledge of the region of the economy outside their local network, we assume that the new opponent is chosen at random from the remaining set of agents.

We explore system behavior when the initial density of the network changes in alternative network evaluation setups and network dynamics scenarios. Preliminary simulation exercises show that, in the non-weighted local MG with *frozen networks*, the smaller the density of the network, the larger average payoffs. If, however, we assume a WMG with *frozen* networks, the population tends to build a network of small clusters composed of highly coordinated agents choosing the same state. Agents of one cluster keep assigning more and more importance to agents of other clusters. These global configurations are not in general robust and can be easily disrupted by subsequent network reassessments. Fluctuating patterns for average payoffs are likely to emerge. Finally, we study what happens when agents play the WMG over *endogenous networks*. In this case, the population splits in two subgroups playing opposite states. Agents in a group are prevalently linked with agents of the other group. Average payoffs converge to one. Thus, the population learns to "globally win" the WMG by selecting the most convenient set of opponents in the game.

The rest of the paper is organized as follows. Section 2 formally describes the model. In Section 3 we present preliminary simulation results. Finally, Section 4 concludes and discusses future developments.

2 The Model

We study a population of myopic agents $i \in I = \{1, ..., N\}$, with N even, playing a minority game in discrete time periods $t = 0, 1, 2, ...$ Each agent i is characterized by its binary state $s_i^t \in \{-1, +1\}$ in the game and by the set $V_i^t \subseteq I - \{i\}$ of other agents is currently linked to (which we call *interaction group*).

We denote by ij_t the directed edge linking agent i to j at time t. Since links are directed, the fact that $j \in V_i^t$ does not necessarily imply that $i \in V_j^t$. The collection $\{V_i^t, i \in I\}$ thus induces at every t a directed graph over I. The set of directed edges (links) held by the agents may be interpreted as their "window over the world". Agents play here a local MG: they observe the part of the population which is relevant to them and adapt to a small and local set of signals.

In each period, agents get a positive (unit) payoff if they play the same as the minority of their interaction group V_i^{t-1} does. More formally, the payoff at time t of agent i is given by:

$$\pi_i^t = \pi_i^t(s \mid n_i^t(s), n_i^t(-s)) = \begin{cases} 1, & \text{if} \quad n_i^t(s) < n_i^t(-s) \\ 0, & \quad \text{otherwise} \end{cases}, \tag{1}$$

where $s\{-1, +1\}$ and $n_i^t(s) = \#\{j \in V_i^{t-1} : s_j^t = s\}$ is the number of agents playing s in V_i^{t-1}.

We assume that, in each time period, any agent is allowed to revise its current state with probability $\theta \in (0, 1]$ and that agents earn the payoff (1) even if they must stay put (or if they do not actually change) their current state. The parameter θ governs the state updating regime. If $\theta = 1$, one has the standard "parallel updating scheme". If θ is sufficiently small, then at most one agent has the option to update its state at t ("asynchronous updating scheme"). Since the nature of the updating scheme has been shown to crucially affect the outcomes of dynamic systems such as the one studied here (Page, 1997), it seems interesting to allow for different updating frequencies in the model (see also Section 4).

Agents are assumed to be myopic and to employ best-response strategies given their local information (Blume, 1995). This means that agent i, when called to revise its state at time t, can only observe time $t-1$ variables concerning agents in V_i^{t-1}. Agents can be considered to remain in a "dormant status" until they are woken-up and allowed to act upon currently available information. Therefore, asynchronous updating ($\theta < 1$) and myopic behavior allow one to model systems where agents behave on the grounds of heterogeneous information.

We consider two alternative setups as to how updating agents evaluate the information coming from their local network and use that information in deciding their next-period state.

In the first setup, every linked agent (that is, every edge ij_{t-1}) always counts as one, independently of the past behavior of the opponent j. Therefore, agent i chooses, with probability θ, its (new) action using a standard, deterministic, best-reply rule:

$$s_i^t = \begin{cases} +1, & \text{if} \quad n_i^{t-1}(+1) < n_i^{t-1}(-1) \\ -1, & \text{if} \quad n_i^{t-1}(+1) > n_i^{t-1}(-1) \end{cases}. \tag{2}$$

We avoid any tie-breaking complications simply by assuming that all interaction groups always contain an odd number of players (see also below). Our results are not dramatically altered if one assumes stochastic TBRs (agents choose at random when a tie occurs) or state-dependent ones (agents stick/switch to their current choice).

In the second setup, we implement a weighted version of the MG (WMG). We suppose that each agent attaches an indicator of importance (*weight*) η_{ij}^t to each link ij_t it maintains (more on how weights are computed and evolve through time below). Given last-period weights η_{ij}^{t-1} attached to agents $j \in V_i^{t-1}$, the agent drawn for state updating – say, i – updates its state by employing a weighted-minority rule as follows:

$$s_i^t = \begin{cases} +1, & \text{if} \quad \sum_{j \in V_i^{t-1}(-1)} \eta_{ij}^{t-1} < \sum_{j \in V_i^{t-1}(+1)} \eta_{ij}^{t-1} \\ \\ -1, & \text{if} \quad \sum_{j \in V_i^{t-1}(-1)} \eta_{ij}^{t-1} > \sum_{j \in V_i^{t-1}(+1)} \eta_{ij}^{t-1} \end{cases}, \tag{3}$$

where we denote by $V_i^{t-1}(s) = \{j \in V_i^{t-1} : s_j^{t-1} = s\}$ and $s \in \{-1, +1\}$. In other words, agents now use a weighted best-reply rule, where opponents do not all count as one and their relative importance is given by their performance in the past.

After state updating (which happens – on average – for at most $\theta \cdot N$ agents in each t), payoffs of *all* agents are computed according to (1). Then, in the WMG setups, weights are updated over time. In each period, after strategies have been updated, each link is assigned a *score* e^t_{ij} equal to 1 or 0, depending on whether the "suggestion" coming from j was correct (that is, if it indicated the state played by the minority of the linked agents) or not. More formally:

$$e^t_{ij} = \begin{cases} 1, & \text{if} \quad \pi^t_i = 1 \quad \text{and} \quad s^t_i \neq s^{t-1}_j \\ 1, & \text{if} \quad \pi^t_i = 0 \quad \text{and} \quad s^t_i = s^{t-1}_j \\ 0, & \text{if} \quad \pi^t_i = 0 \quad \text{and} \quad s^t_i \neq s^{t-1}_j \\ 0, & \text{if} \quad \pi^t_i = 1 \quad \text{and} \quad s^t_i = s^{t-1}_j \end{cases} . \tag{4}$$

Agents then use the score to update the weight of each link as follows:

$$\eta^t_{ij} = \alpha\eta^{t-1}_{ij} + (1-\alpha)e^t_{ij}, \tag{5}$$

where $\alpha \in [0, 1]$ and $\eta^t_{ij} \in [0, 1]$.

Notice that by applying the weight updating rule in (5), agents trade-off their memory about the "past contributions" from agent j with its "current contribution" (e^t_{ij}). The parameter α tunes the memory effect. If $\alpha \simeq 0$, weights track very closely current scores (no memory), whereas, if $\alpha \simeq 1$, memory becomes very important. The η^t_{ij} series are smoother and are quite robust to new scores. See Figure 1 for an example of weight series generated by applying rule (5) to the same series of e^t_{ij} for alternative values of α (cf. Skyrms and Pemantle (2000), Weisbuch, Kirman, and Herreiner (2000) and Kirman and Vriend (2001) for a similar approach).

Weights provide an indication of how useful links have been in the past in helping to take a correct decision, no matter if the player actually did choose a new state or was forced to stick to its current one. In a sense, the collection $\{\eta^t_{ij}, j \in V^{t-1}_i\}$ is a measure of how much i's interaction group has been able to effectively and consistently convey information to agent i. On average, the closer all weights $\{\eta^t_{ij}, j \in V^{t-1}_i\}$ to 1, the more agent i can rest on the suggestions coming from its local network in taking a state updating decision

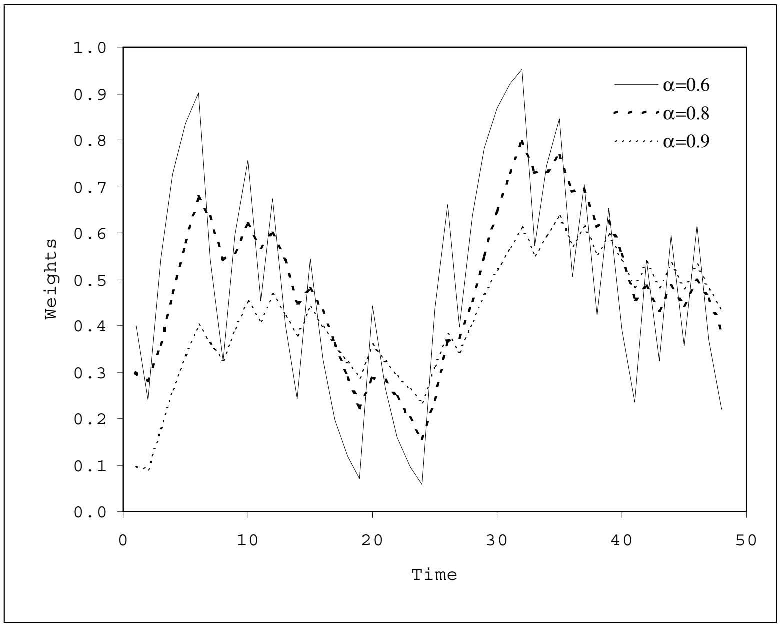

Figure 1: An example of the dynamics of individual weights. Equation (5) is applied over the same time-series of scores e^t_{ij} for three values of the memory parameter α =0.6, 0.8, 0.9.

at time t.

Therefore, the index

$$H_i^t = \frac{\sum_{j \in V_i^{t-1}} \left(\eta_{ij}^t\right)^2}{|V_i^{t-1}|} \in [0,1] \tag{6}$$

provides information about the overall helpfulness of agent i's interaction group V_i^{t-1}. We employ squared weights to emphasize the contribution to H_i^t made by extreme-valued links, that is opponents which are considered either as completely useless or completely helpful.

Finally, let us describe how networks evolve through time. At time $t = 0$, each agent is randomly linked with L^0 other agents. More formally, each agent is endowed with the same (odd) number of links $L_i^0 = L^0 = \lfloor \delta N \rfloor$, where $\delta \in (0,1)$ is a proxy for the "density" of the initial network and $L^0 \in \{1, 3, 5, ..., N-1\}$.

We experiment with two alternative scenarios as far as network dynamics is concerned. In the first one, *frozen networks* are assumed. Links cannot change in the entire process and, under the WMG setup, weights can only be updated. This allows one to disentangle the roles played by coordination, local interactions, and, possibly, weight updating in shaping collective behavior.

In the second setup, we study an *endogenous networks* system. After state updating, agents are allowed to discard badly-performing links. More formally, we introduce a threshold $\mu \in (0,1)$ and we suppose that agent i deletes all links such that:

$$\eta_{ij}^t < \mu \frac{\sum_{h \in V_i^{t-1}} \eta_{ih}^t}{|V_i^{t-1}|} = \mu \bar{\eta}_i^t. \tag{7}$$

In line with the local nature of information diffusion in our economy, we assume that agents replace any discarded link with a new one (randomly chosen from the set of all currently unconnected agents) and that any new link is initially assigned a weight equal to the average of the weights of all undiscarded links. Thus, the number of links that each agent holds remains constant at any t (equal to L^0).

3 Simulation Results

In this Section, we present some preliminary Monte-Carlo simulation exercises and discuss the most important properties displayed by the collective behavior of the system.

We begin with a benchmark parametrization where: (i) the frequency of state updating is $\theta = 0.20$; (ii) the population size is $N = 100$.

Our main goal is to explore how networks influence collective dynamics. Therefore, we start by studying what happens to the distribution of individual payoffs and to the interaction structure (i.e. links and weights) when the "density" of the initial network (δ or, equivalently, L^0) changes in each of the following three setups:

1. Agents play a *non-weighted* MG and networks are *frozen*.

2. Agents play a *weighted* MG and networks are *frozen*.

3. Agents play a *weighted* MG and networks are *endogenous*.

In all our exercises, we average our observations over $M = 10$ independent Monte-Carlo simulations in order to wash away across-simulation variability. Moreover, we observe the system dynamics until the economy has reached a sufficiently stable behavior, which typically happens for $2000 \leq T \leq 4000$ for all (α, θ, N) setups explored in our simulations.

All foregoing results are quite robust to alternative (α, θ) parameterizations, as well as to larger Montecarlo sample sizes (M), population sizes (N) and time horizons (T). See, however, Section 4 for a brief discussion on the need for a deeper sensitivity analysis over the whole parameter space.

3.1 Non-Weighted Minority Games over Frozen Networks

In the first set of simulations, we study a system where agents play a non-weighted MG over exogenously fixed networks. This means that the initial interaction structure $\{V_i^0, i \in I\}$ is not allowed to change. Agents always observe the same local network and are rewarded only if they play as the minority of their opponents. Furthermore, each opponent always counts as one, irrespective of its past behavior – weights are always equal to one for each link.

In such a setting, an interesting question arises about whether collective coordination can be affected by the size of the observation window held by the agents. In Figure 2 we plot the evolution over time of population-average payoffs

$$\overline{\pi}^t = \frac{1}{N} \sum_{i=1}^{N} \pi_i^t \tag{8}$$

as the number of links with which we endow each agent (L^0) changes.

Populations with a smaller number of links provide, on average, higher payoff levels. A better coordination is then achieved by playing with a smaller number of opponents. In fact, if agents are only able to observe the world through a smaller window, it is more likely that these windows do not overlap (that is, $i \in V_j^0$ but $j \notin V_i^0$). This allows agents better to coordinate and the system to reach a good average performance.

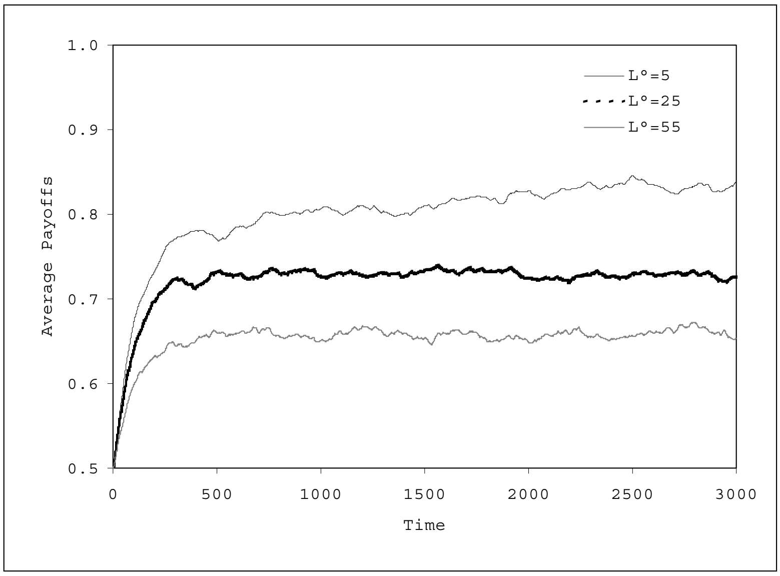

Figure 2: Montecarlo mean of population-average payoffs when agents play a *non-weighted* MG over *frozen* networks and the initial number of links changes (L^0=5, 25, 55). Montecarlo mean performed over $M = 10$ runs. Setup: $N = 100$, $\theta = 0.20$.

Notice also that the high instability and frustrated behaviors often displayed in standard MG are avoided (for similar findings in the context of nearest-neighbors interactions on lattices, cf. Burgos, Ceva, and Perazzo (2004), Kalinowski, Schulz, and Briese (2000) and Moelbert and De Los Rios (2002)). Our results suggest that local interactions can partly replace individual memory: a stable collective behavior characterized by a better-than-average coordination can be reached even if players cannot observe the last global winning sides and cannot learn in the strategy space (cf. also Bottazzi, Devetag, and Dosi (2003)).

3.2 Weighted Minority Games over Frozen Networks

Let us turn now to the case where networks are still frozen but agents can weigh the importance of the links they maintain. Notice that, in this setup, agents cannot discard poorly performing links to add others, but they can attach different beliefs about whether the information coming to a particular opponent is useful in deciding their next-period state.

Each opponent j in a given V_i^0 then contributes to the decision at time t in proportion to its link-weight η_{ij}^t. In what follows, we suppose that the importance of memory in weights dynamics is very high ($\alpha = 0.99$). This means that the score e_{ij}^t currently earned by the opponent j has a small impact in changing the assessments of agent i about its beliefs over j.

When agents are allowed to "select which details to observe from their window", the collective behavior of the system sensibly changes as compared to the non-weighted case. Indeed, populations holding a large number of links are now able to focus on smaller subsets of trusted connections. The latter are the ones that, in the past, have guaranteed high scores (see eq. 4). As a consequence, the population tends to break down into groups of agents sharing the same state, which, however maintain very weak (if possibly no) connections among them. Agents belonging to, say, a $+1$ group instead build connections with players belonging to a -1 group.

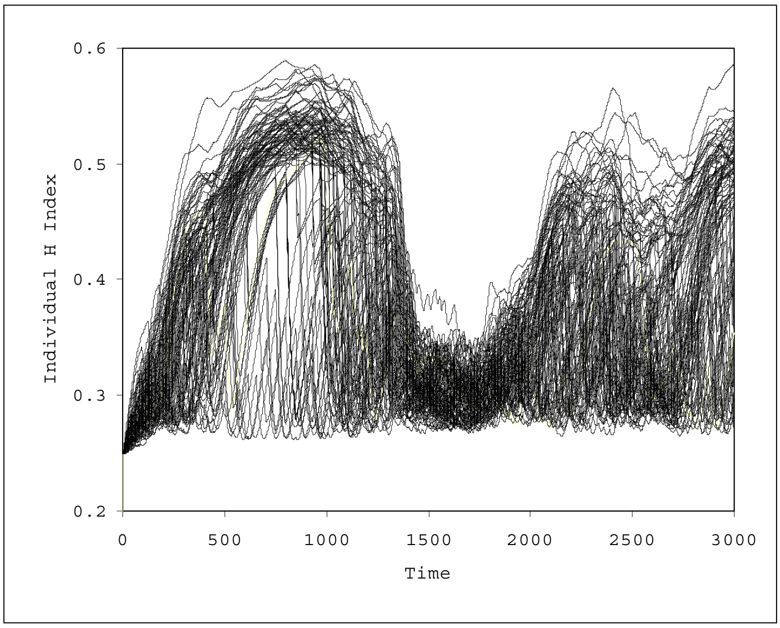

Figure 3: Time-series of individual H_i^t indices, $i = 1, ..., N$, when agents play a *weighted* MG over *frozen* networks. Values of H_i^t refer to a benchmark run. Setup: $N = 100$, $\alpha = 0.99$, $\theta = 0.20$, $L^0 = 55$.

Such a clustering process entails a high degree of coordination among agents: it is therefore very difficult to achieve and very sensitive to small fluctuations. To see this, we report in Figure 3 all individual indices H_i^t (see eq. 6) for a typical simulation run. Notice

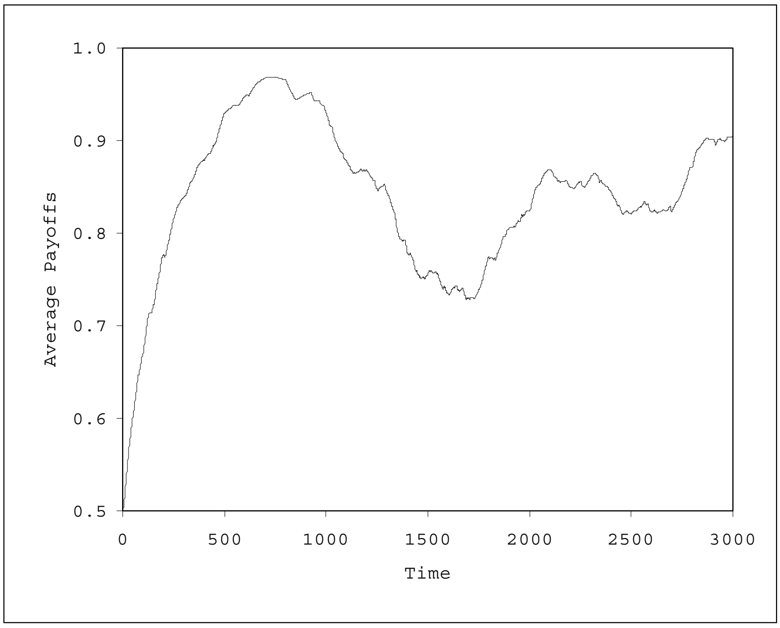

Figure 4: Time-series of population average payoffs when agents play a *weighted* MG over *frozen* networks. Payoff values refer to the same benchmark run of Figure 3. Setup: $N = 100$, $\alpha = 0.99$, $\theta = 0.20$, $L^0 = 55$.

that all indices increase when a feasible coordination pattern is under formation. However, after such a clustered network has been built, any small change in the global state-weight configuration is able to disrupt the coordination pattern that has just emerged (for a similar evidence in the context of open-ended economies self-organizing "on the edge of chaos", cf. Lindgren (1991)).

Population-average payoffs accordingly increase during the construction of the clustered network and then falls after the latter has been destroyed, see Figure 4. The initial wave (with H_i^t indices all increasing) is due to a larger and larger number of agents strengthening their links to agents belonging to an opposite-state cluster. While more and more agents join the cluster, the network falls apart. This causes a huge drop in both payoffs and H_i^t values.

This oscillating behavior is not limited to the early stages of the simulation, but can be observed at any time-scale. In a further set of exercises not reported here (but available to the reader on request), we observed continuing oscillations until $t = 50000$, and, more generally, well beyond the expected relaxation horizon.

Interestingly enough, populations with an initially low number of links get now *lower* average payoffs than their more connected counterparts (see Figure 5). This is because weight assessment over a small number of agents does not provide a large enough pool from which selecting a robust subset of opponents behaving in a consistently useful way over time.

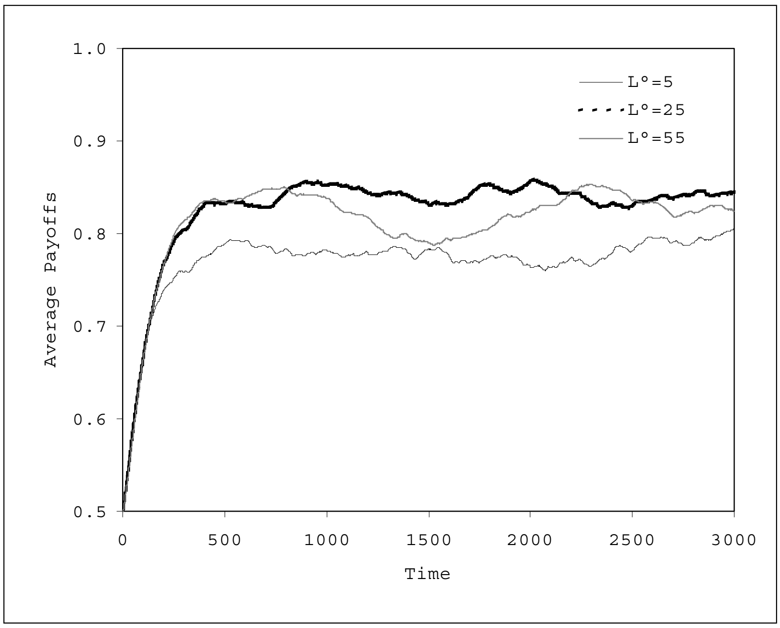

Figure 5: Montecarlo mean of population-average payoffs when agents play a *weighted* MG over *frozen* networks and the initial number of links changes (L^0=5, 25, 55). Montecarlo mean performed over $M = 10$ runs. Setup: $N = 100$, $\alpha = 0.99$, $\theta = 0.20$.

Note, however, that all three populations reach now (on average) higher payoff levels than in the previous case. This suggests that allowing for some endogeneity in network formation (e.g., payoff-dependent weight dynamics) implies a better collective performance. To explore this intuition further, we move now to the *endogenous networks* setup.

3.3 Weighted Minority Games over Endogenous Networks

Let us now suppose that agents play a WMG and can endogenously delete badly-performing links. We set the cutoff value to $\mu = 0.50$. This means that an agent holding a link ij_t such that $\eta_{ij}^t < \frac{1}{2}\overline{\eta}_i^t$ will delete it and replace it by another ij_t', where j' is drawn at random from the pool of all currently non-linked players.

In this setup, the coordination process becomes quite efficient and very rapid. Agents typically split into two (almost) equally-sized groups: players in the first persistently choose the state +1, those in the other always select the state −1. Agents in one group maintain a large majority of links with agents in the other group. This allows them effortlessly to get a positive reward and rapidly converge to an overall system configuration characterized by average payoffs $\overline{\pi}^t$ very close (if not equal) to one. Therefore, despite the local nature of interaction patterns, the population is able to "globally win" the MG by endogenously selecting the right set of opponents.

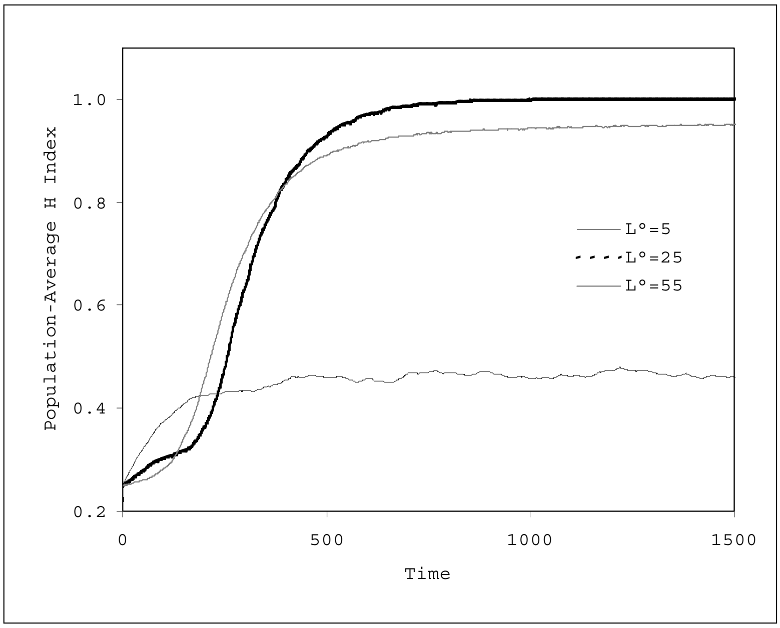

Figure 6: Montecarlo mean of population average $\overline{H}^t$ index when agents play a *weighted* MG over *endogenous* networks and the initial number of links changes (L^0=5, 25, 55). The cutoff value for all three populations is $\mu = 0.50$. Montecarlo means performed over $M = 10$ runs. Setup: $N = 100$, $\alpha = 0.99$, $\theta = 0.20$.

What is more, convergence to a stable interaction pattern allowing for an efficient collective behavior does not generally depend on the initial degree of connectivity. Figure 6 shows the population-average of H_i^t indices:

$$\overline{H}^t = \frac{1}{N}\sum_{i=1}^{N} H_i^t, \tag{9}$$

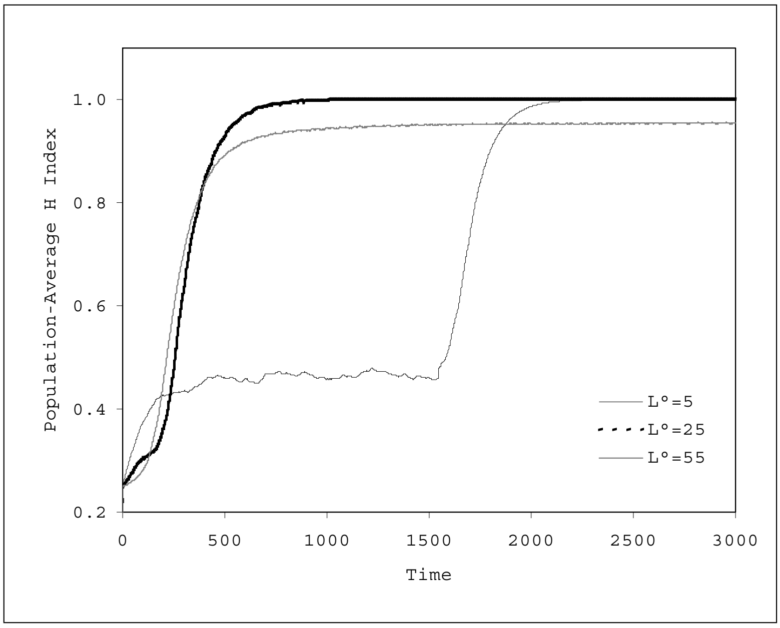

Figure 7: Continuation of the simulation runs presented in Figure 6 after the cutoff value for the $L^0 = 5$ population has been increased to $\mu = 0.80$ at $t = 1550$. Montecarlo mean of population average $\overline{H}^t$ index performed over $M = 10$ runs. Setup: $N = 100$, $\alpha = 0.99$, $\theta = 0.20$.

for three populations characterized by $L^0 = 5, 25, 55$. The two *initially* highly connected populations (i.e. $L^0 = 25, 55$) quickly converge to a steady-state with very high $\overline{H}^t$ values. Average payoffs $\overline{\pi}^t$ for these two populations quickly converge towards 1.

Incidentally, notice that the average $\overline{H}^t$ index follows a s-shaped pattern, i.e. the signature of a diffusion process (Dosi, 1991). Indeed, when the two groups start forming, an agent which is not yet part of any group can simply join one of them by selecting a linked agent in the other group, similarly to what happens in an epidemic process.

Only the $L^0 = 5$ population fails to stabilize. The H_i^t values are rather small, as well as the population-average of weights. This prevents the worst performing links from falling below the deletion threshold. Therefore, agents cannot effectively employ network adaptation to reach coordination.

However, if the cutoff value increases, agents find it easier to get rid of badly-performing links. To see whether this is sufficient to trigger convergence to the efficient state, we raised on-the-fly (around $t = 1550$) the cutoff values of agents belonging to the $L^0 = 5$ population to $\mu = 0.80$. Figure 7 shows the continuation of the simulation presented in Figure 6. Now the initially weakly-connected population is also able to converge to the efficient stable state.

4 Concluding Remarks

In this paper, we study a local version of the MG, where agents initially hold directed edges connecting them with other players in the population. Agents can only observe the state chosen by their opponents in the last-period and care about being in the minority of their own interaction group. To choose their next-period state, agents employ simple, best-reply rules.

We have explored two network evaluation setups. In the first, links are non-weighted and agents simply count the number of their opponents in either state to decide which side to take in the next period. In the second, links (that is, opponents) are attached a *weight*. A weight evolves in a path-dependent way through time: its value increases only if the information provided by the linked agent has been helpful in the past.

We start from a scenario where networks are *frozen* (links cannot be deleted/added) and we then move to a system where *endogenous* networks are assumed (agents can discard badly-performing links and replace them with other ones chosen at random).

Our results indicate that the very possibility of playing the MG locally and endogenously acting over the network structure might strongly affect the efficiency of collective behavior. For example, even when agents cannot display path-dependent discrimination among connections and must stick to an exogenously given network, efficiency can be increased if agents hold small, local interaction groups. In a sense, individual memory can be substituted for the information locally gathered by the players in a myopic way.

Furthermore, simulations show that the efficiency of the system can be greatly enhanced by allowing players to act upon the structure of the network. Indeed, our results suggest that populations playing a WMG over frozen networks are able to self-organize and build transient clustered networks that attain high payoff levels. These self-organized configurations are very sensitive to subsequent network reassessments: their disruption may generate abrupt fluctuations in average payoffs.

Such huge fluctuations, as well as sub-optimal payoff levels, may be avoided, however, if agents are able to delete badly performing links and replace them with random ones. Then, the population is able to "globally win" the WMG and consistently reach a stable state where all agents get a positive payoff. Provided that agents are sufficiently willing to discard links, this conclusion holds independently of the initial number of connections assigned to each player.

These promising results should be more carefully scrutinized *vis-à-vis* an extensive Monte-Carlo analysis over the entire parameter space. For example, the role of the frequency of state updating (θ) in shaping long-run properties of the system should be addressed in more detail.

Furthermore, we have shown that an adequate level for the cut-off parameter μ is able to quickly move the whole population towards an efficient state. Although we were interested here in studying the very possibility of this event, it might be interesting to investigate more deeply how this threshold value varies with N and L^0. Preliminary exercises show that the cut-off level required to trigger the replacement of badly performing links seems to depend on the absolute number of links L^0 only, rather than on the ratio L^0/N, as one could expect. The reason why this happens could be traced back to the initial distribution of link weights $\{\eta_{ij}^0\}$. In fact, our link replacement process possesses a self-reinforcing nature: if agent i replaces a link at time t, its average evaluation $\overline{\eta}_i^t$ increases. This makes more likely that agent i will replace another link at $t+1$. This process goes on until all links have very similar evaluations. However, to start the replacement process in the first place, we need at least one link falling sufficiently below the average level, and this may depend on the distribution of evaluations of the initial, randomly drawn evaluations over the L^0 links.

Finally, the relationships between network structure and collective behavior should be studied more deeply. For example, it might be interesting to assess how the fine properties of the network structure initially in place (e.g., presence and number of cycles, structural and locational properties, etc.) affect the coupled network-state dynamics in the system (as well as the emergent network structure). More generally, the robustness of our results might be tested against alternative network formation processes allowing for score functions different from (4) and endogenously changing interaction group sizes.

References

Anghel, M., Z. Toroczkai, K. Bassler and G. Korniss (2004), "Competition-driven network dynamics: emergence of a scale-free leadership structure and collective efficiency.", *Physical Review Letters*, 92: 058701.

Arthur, W. (1994), "Inductive Reasoning and Bounded Rationality", *American Economic Review, Papers and Proceedings*, 84: 406–411.

Blume, L. (1995), "The statistical mechanics of best-response strategy revision", *Games and Economic Behavior*, 11: 111–145.

Bottazzi, G., G. Devetag and G. Dosi (2003), "Adaptive Learning and Emergent Coordination in Minority Games", in Cowan, R. and N. Jonard (eds.), *Heterogeneous Agents, Interactions and Economic Performance*. Lecture Notes in Economics and Mathematical Systems, Berlin, Springer Verlag.

Burgos, E., H. Ceva and R. Perazzo (2004), "The Evolutionary Minority Game with Local Coordination", *Physica A*, 337: 635–644.

Challet, D., M. Marsili and Y. Zhang (2000), "Modeling Market Mechanism with Minority Game", *Physica A*, 276: 284–315.

Challet, D. and Y. Zhang (1997), "Emergence of cooperation and organization in an evolutionary game", *Physica A*, 226: 407–418.

Challet, D. and Y. Zhang (1998), "On the Minority Game: Analytical and Numerical Studies", *Physica A*, 256: 514–532.

Dosi, G. (1991), "The research on innovation diffusion: an assessment", in Grubler, A. and N. Nakicenovic (eds.), *Diffusion of technologies and social behavior*. New York, Springer Verlag.

Fagiolo, G. (1998), "Spatial interactions in dynamic decentralized economies: a review", in Cohendet, P., P. Llerena, H. Stahn and G. Umbhauer (eds.), *The Economics of Networks. Interaction and Behaviours*. Berlin - Heidelberg, Springer Verlag.

Fagiolo, G. (2005), "Endogenous Neighborhood Formation in a Local Coordination Model with Negative Network Externalities", *Journal of Economic Dynamics and Control*, 29: 297–319.

Fagiolo, G., L. Marengo and M. Valente (2004), "Endogenous Networks in Random Population Games", *Mathematical Population Studies*, 11: 121–147.

Kalinowski, T., H. Schulz and M. Briese (2000), "Cooperation in the Minority Game with local information", *Physica A*, 277: 502–508.

Kauffman, S. (1993), *The origins of order*. Oxford, Oxford University Press.

Kirley, M. (2004), "Evolutionary minority games with small-world interactions", Department of Computer Science and Software Engineering, The University of Melbourne, Victoria, Australia.

Kirman, A. (1997), "The Economy as an Evolving Network", *Journal of Evolutionary Economics*, 7: 339–353.

Kirman, A. and N. Vriend (2001), "Evolving Market Structure: An ACE Model of Price Dispersion and Loyalty", *Journal of Economic Dynamics and Control*, 25: 459–502.

Lindgren, K. (1991), "Evolutionary phenomena in simple dynamics", in Langton, C., C. Taylor, J. Farmer and S. Rasmussen (eds.), *Artificial Life II*. Santa Fe Institute, Santa Fe and Reading, MA, Addison-Wesley.

Lo, T., K. Chan, P. Hui and N. Johnson (2004), "Theory of enhanced performance emerging in a sparsely-connected competitive population", cond-mat/0409140.

Marsili, M. and D. Challet (2001), "Trading Behavior and Excess Volatility in Toy Markets", *Advances in Complex Systems*, 4: 3–17.

Moelbert, S. and P. De Los Rios (2002), "The Local Minority Game", *Physica A*, 303: 217–227.

Ochs, J. (1990), "The Coordination Problem in the Centralized Markets: An Experiment", *Quarterly Journal of Economics*, 105: 545–559.

Ochs, J. (1995), "Coordination Problems", in Kagel, J. and A. Roth (eds.), *Handbook of Experimental Economics*. Princeton, Princeton University Press.

Paczuski, M., K. Bassler and A. Corral (2000), "Self-organized Networks of Competing Boolean Agents", *Physical Review Letters*, 84: 3185–3188.

Page, S. (1997), "On Incentives and Updating in Agent Based Models", *Computational Economics*, 10: 67–87.

Skyrms, B. and R. Pemantle (2000), "A dynamic model of social network formation", *Proceedings of the National Academy of Sciences*, 97: 9340–9346.

Slanina, F. (2000), “Social organization in the Minority Game model”, *Physica A*, 286: 367–376.

Weisbuch, G., A. Kirman and D. Herreiner (2000), “Market Organisation and Trading Relationships”, *Economic Journal*, 110: 411–436.